



OPEN

Online speech synthesis using a chronically implanted brain–computer interface in an individual with ALS

Miguel Angrick^{1✉}, Shiyu Luo², Qinwan Rabbani³, Daniel N. Candrea², Samyak Shah¹, Griffin W. Milsap⁴, William S. Anderson⁵, Chad R. Gordon^{5,6}, Kathryn R. Rosenblatt^{1,7}, Lora Clawson¹, Donna C. Tippett^{1,8,9}, Nicholas Maragakis¹, Francesco V. Tenore⁴, Matthew S. Fifer⁴, Hynek Hermansky^{10,11}, Nick F. Ramsey¹² & Nathan E. Crone^{1✉}

Brain–computer interfaces (BCIs) that reconstruct and synthesize speech using brain activity recorded with intracranial electrodes may pave the way toward novel communication interfaces for people who have lost their ability to speak, or who are at high risk of losing this ability, due to neurological disorders. Here, we report online synthesis of intelligible words using a chronically implanted brain–computer interface (BCI) in a man with impaired articulation due to ALS, participating in a clinical trial (ClinicalTrials.gov, NCT03567213) exploring different strategies for BCI communication. The 3-stage approach reported here relies on recurrent neural networks to identify, decode and synthesize speech from electrocorticographic (ECoG) signals acquired across motor, premotor and somatosensory cortices. We demonstrate a reliable BCI that synthesizes commands freely chosen and spoken by the participant from a vocabulary of 6 keywords previously used for decoding commands to control a communication board. Evaluation of the intelligibility of the synthesized speech indicates that 80% of the words can be correctly recognized by human listeners. Our results show that a speech-impaired individual with ALS can use a chronically implanted BCI to reliably produce synthesized words while preserving the participant’s voice profile, and provide further evidence for the stability of ECoG for speech-based BCIs.

A variety of neurological disorders, including amyotrophic lateral sclerosis (ALS), can severely affect speech production and other purposeful movements while sparing cognition. This can result in varying degrees of communication impairments, including Locked-In Syndrome (LIS)^{1,2}, in which patients can only answer yes/no questions or select from sequentially presented options using eyeblinks, eye movements, or other residual movements. Individuals such as these may use augmentative and alternative technologies (AAT) to select among options on a communication board, but this communication can be slow, effortful, and may require caregiver intervention. Recent advances in implantable brain–computer interfaces (BCIs) have demonstrated the feasibility of establishing and maintaining communication using a variety of direct brain control strategies that bypass

¹Department of Neurology, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ²Department of Biomedical Engineering, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ³Department of Electrical and Computer Engineering, The Johns Hopkins University, Baltimore, MD, USA. ⁴Research and Exploratory Development Department, Johns Hopkins Applied Physics Laboratory, Laurel, MD, USA. ⁵Department of Neurosurgery, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ⁶Section of Neuroplastic and Reconstructive Surgery, Department of Plastic Surgery, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ⁷Department of Anesthesiology & Critical Care Medicine, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ⁸Department of Otolaryngology-Head and Neck Surgery, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ⁹Department of Physical Medicine and Rehabilitation, The Johns Hopkins University School of Medicine, Baltimore, MD, USA. ¹⁰Center for Language and Speech Processing, The Johns Hopkins University, Baltimore, MD, USA. ¹¹Human Language Technology Center of Excellence, The Johns Hopkins University, Baltimore, MD, USA. ¹²UMC Utrecht Brain Center, Department of Neurology and Neurosurgery, University Medical Center Utrecht, Utrecht, The Netherlands. ✉email: mangric1@jhu.edu; ncrone@jhmi.edu

weak muscles, for example to control a switch scanner^{3,4}, a computer cursor⁵, to write letters⁶ or to spell words using a hybrid approach of eye-tracking and attempted movement detection⁷. However, these communication modalities are still slower, more effortful, and less intuitive than speech-based BCI control⁸.

Recent studies have also explored the feasibility of decoding attempted speech from brain activity, outputting text or even acoustic speech, which could potentially carry more linguistic information such as intonation and prosody. Previous studies have reconstructed acoustic speech in offline analysis from linear regression models⁹, convolutional¹⁰ and recurrent neural networks^{11,12}, and encoder-decoder architectures¹³. Concatenative approaches from the text-to-speech synthesis domain have also been explored^{14,15}, and voice activity has been identified in electrocorticographic (ECoG)¹⁶ and stereotactic EEG recordings¹⁷. Moreover, speech decoding has been performed at the level of American English phonemes¹⁸, spoken vowels^{19,20}, spoken words²¹ and articulatory gestures^{22,23}.

Until now, brain-to-speech decoding has primarily been reported in individuals with unimpaired speech, such as patients temporarily implanted with intracranial electrodes for epilepsy surgery. To date, it is unclear to what extent these findings will ultimately translate to individuals with motor speech impairments, as in ALS and other neurological disorders. Recent studies have demonstrated how neural activity acquired from an ECoG grid²⁴ or from microelectrodes²⁵ can be used to recover text from a patient with anarthria due to a brainstem stroke, or from a patient with dysarthria due to ALS, respectively. Prior to these studies, a landmark study allowed a locked-in volunteer to control a real-time synthesizer generating vowel sounds²⁶. More recently, Metzger et al.²⁷ demonstrated in a clinical trial participant diagnosed with quadriplegia and anarthria a multimodal speech-neuroprosthetic system that was capable of synthesizing sentences in a cued setting from silent speech attempts. In our prior work, we presented a ‘plug-and-play’ system that allowed a clinical trial participant living with ALS to issue commands to external devices, such as a communication board, by using speech as a control mechanism²⁸.

In related work, BCIs based on non-invasive modalities, such as electroencephalography (EEG), functional near-infrared spectroscopy (fNIRS) or functional magnetic resonance imaging (fMRI) have been investigated for speech decoding applications. These studies have largely focused on imagined speech²⁹ to avoid contamination by movement artifacts³⁰. Recent work by Dash et al., for example, reported speech decoding results for imagined and spoken phrases from 3 ALS patients using magnetoencephalography (MEG)³¹. While speech decoding based on non-invasive methodologies is an important branch in the BCI field as they do not require a surgery and may be adopted by a larger population more easily, their current state of the art comes with disadvantages compared to implantable BCIs as they lack either temporal or spatial resolution, or are currently not feasible for being used at home.

Here, we show that an individual living with ALS and participating in a clinical trial of an implantable BCI (ClinicalTrials.gov, NCT03567213) was able to produce audible, intelligible words that closely resembled his own voice, spoken at his own pace. Speech synthesis was accomplished through online decoding of ECoG signals generated during overt speech production from cortical regions previously shown to represent articulation and phonation, following similar previous work^{11,19,32,33}. Our participant had considerable impairments in articulation and phonation. He was still able to produce some words that were intelligible when spoken in isolation, but his sentences were often unintelligible. Here, we focused on a closed vocabulary of 6 keywords, originally used for decoding spoken commands to control a communication board. Our participant was capable of producing these 6 keywords individually with a high degree of intelligibility. We acquired training data over a period of 6 weeks and deployed the speech synthesis BCI in several separate closed-loop sessions. Since the participant could still produce speech, we were able to easily and reliably time-align the individual’s neural and acoustic signals to enable a mapping between his cortical activity during overt speech production processes and his voice’s acoustic features. We chose to provide delayed rather than simultaneous auditory feedback in anticipation of ongoing deterioration in the patient’s speech due to ALS, with increasing discordance and interference between actual and BCI-synthesized speech. This design choice would be ideal for a neuroprosthetic device that remains capable of producing intelligible words as an individual’s speech becomes increasingly unintelligible, as was expected in our participant due to ALS.

Here, we present a self-paced BCI that translates brain activity directly to acoustic speech that resembles characteristics of the user’s voice profile, with most synthesized words of sufficient intelligibility to be correctly recognized by human listeners. This work makes an important step in adding more evidence that recent speech synthesis from neural signals in patients with intact speech can be translated to individuals with neurological speech impairments, by first focusing on a closed vocabulary that the participant can reliably generate at his own pace, before generalizing towards unseen words. Synthesizing speech from the neural activity associated with overt speech allowed us to demonstrate the feasibility of reproducing the acoustic features of speech when ground truth is available and its alignment with an acoustic target is straightforward, in turn setting a standard for future efforts when ground truth is unavailable, as in the Locked In Syndrome. Moreover, because our speech synthesis model was trained on data that preceded testing by several months, our results also support the stability of ECoG as a basis for speech BCIs.

Approach

In order to synthesize acoustic speech from neural signals, we designed a pipeline that consisted of three recurrent neural networks (RNNs) to (1) identify and buffer speech-related neural activity, (2) transform sequences of speech-related neural activity into an intermediate acoustic representation, and (3) eventually recover the acoustic waveform using a vocoder. Figure 1 shows a schematic overview of our approach. We acquired ECoG signals from two electrode grids that covered cortical representations for speech production including ventral sensorimotor cortex and the dorsal laryngeal area (Fig. 1A). Here, we focused only on a subset of electrodes that had previously been identified as showing significant changes in high-gamma activity associated with overt

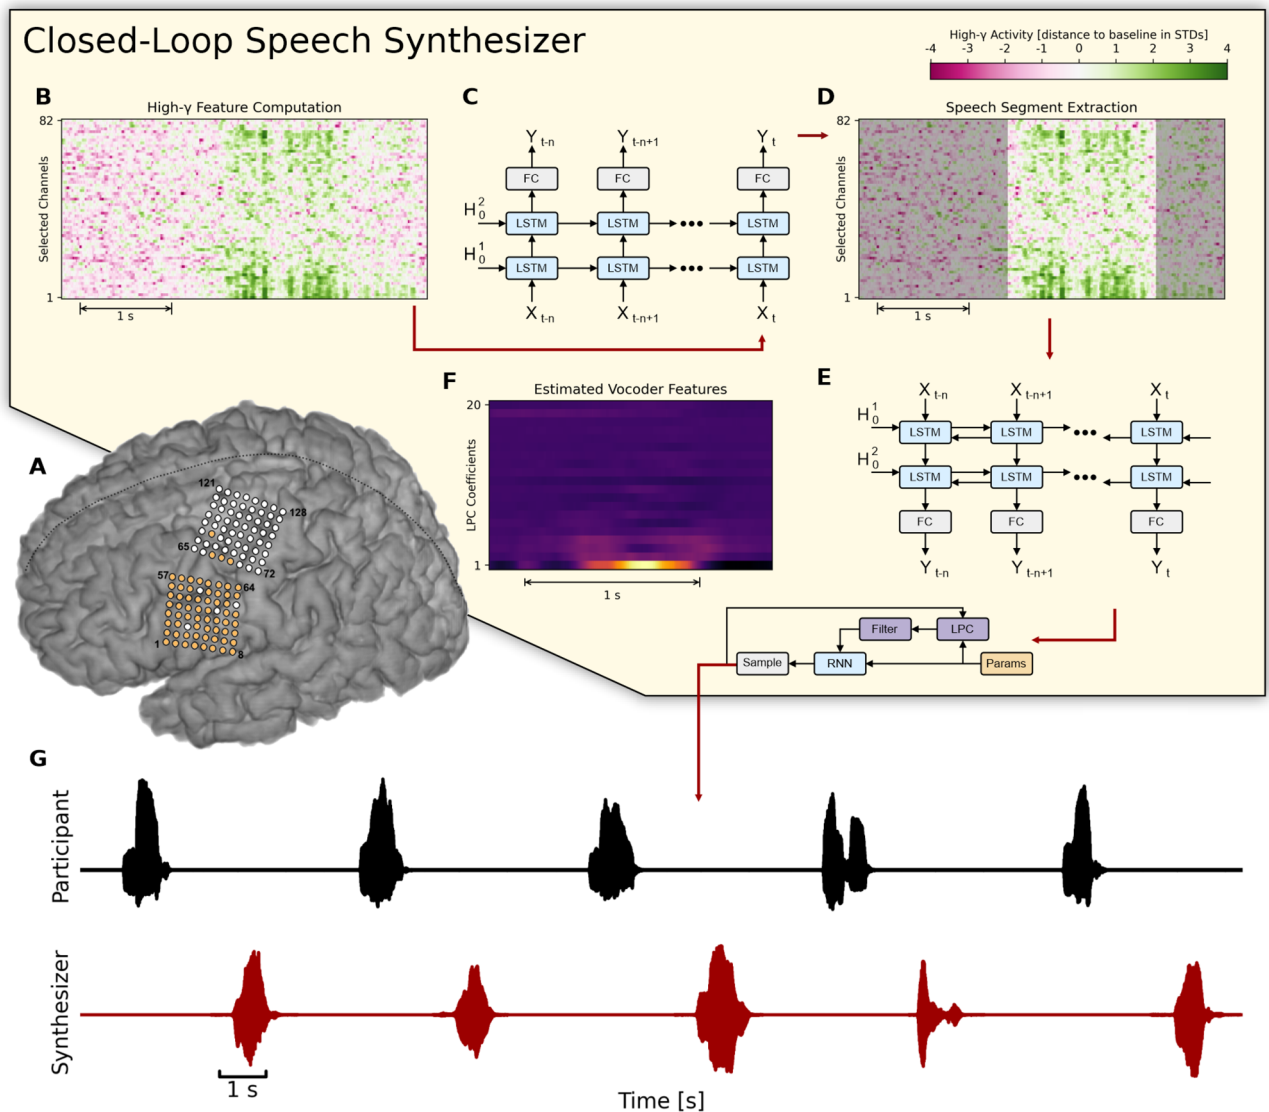


Figure 1. Overview of the closed-loop speech synthesizer. (A) Neural activity is acquired from a subset of 64 electrodes (highlighted in orange) from two 8×8 ECoG electrode arrays covering sensorimotor areas for face and tongue, and for upper limb regions. (B) The closed-loop speech synthesizer extracts high-gamma features to reveal speech-related neural correlates of attempted speech production and propagates each frame to a neural voice activity detection (nVAD) model (C) that identifies and extracts speech segments (D). When the participant finishes speaking a word, the nVAD model forwards the high-gamma activity of the whole extracted sequence to a bidirectional decoding model (E) which estimates acoustic features (F) that can be transformed into an acoustic speech signal. (G) The synthesized speech is played back as acoustic feedback.

speech production (see Supplementary Fig. 2). From the raw ECoG signals, our closed-loop speech synthesizer extracted broadband high-gamma power features (70–170 Hz) that had previously been demonstrated to encode speech-related information useful for decoding speech (Fig. 1B)^{10,14}.

We used a unidirectional RNN to identify and buffer sequences of high-gamma activity frames and extract speech segments (Fig. 1C,D). This neural voice activity detection (nVAD) model internally employed a strategy to correct misclassified frames based on each frame's temporal context, and additionally included a context window of 0.5 s to allow for smoother transitions between speech and non-speech frames. Each buffered sequence was forwarded to a bidirectional decoding model that mapped high-gamma features onto 18 Bark-scale cepstral coefficients³⁴ and 2 pitch parameters, henceforth referred to as LPC coefficients^{35,36} (Fig. 1E,F). We used a bidirectional architecture to include past and future information while making frame-wise predictions. Estimated LPC coefficients were transformed into an acoustic speech signal using the LPCNet vocoder³⁶ and played back as delayed auditory feedback (Fig. 1G).

Results

Synthesis performance

When deployed in sessions with the participant for online decoding, our speech-synthesis BCI was reliably capable of producing acoustic speech that captured many details and characteristics of the voice and pacing of the participant’s natural speech, often yielding a close resemblance to the words spoken in isolation from the participant. Figure 2A provides examples of original and synthesized waveforms for a representative selection of words time-aligned by subtracting the duration of the extracted speech segment from the nVAD. Onset timings from the reconstructed waveforms indicate that the decoding model captured the flow of the spoken word while also synthesizing silence around utterances for smoother transitions. A comparison between voice activity for spoken and synthesized speech revealed a median Levenstein distance of 235 ms, hinting that the synthesis approach was capable of generating speech that adequately matched the timing of the spoken counterpart. Figure 2B shows the corresponding acoustic spectrograms for the spoken and synthesized words, respectively. The spectral structures of the original and synthesized speech shared many common characteristics and achieved average correlation scores of 0.67 (± 0.18 standard deviation) suggesting that phoneme and formant-specific information were preserved.

We conducted 3 sessions across 3 different days (approximately 5 and a half months after the training data was acquired, each session lasted 6 min) to repeat the experiment with acoustic feedback from the BCI to the

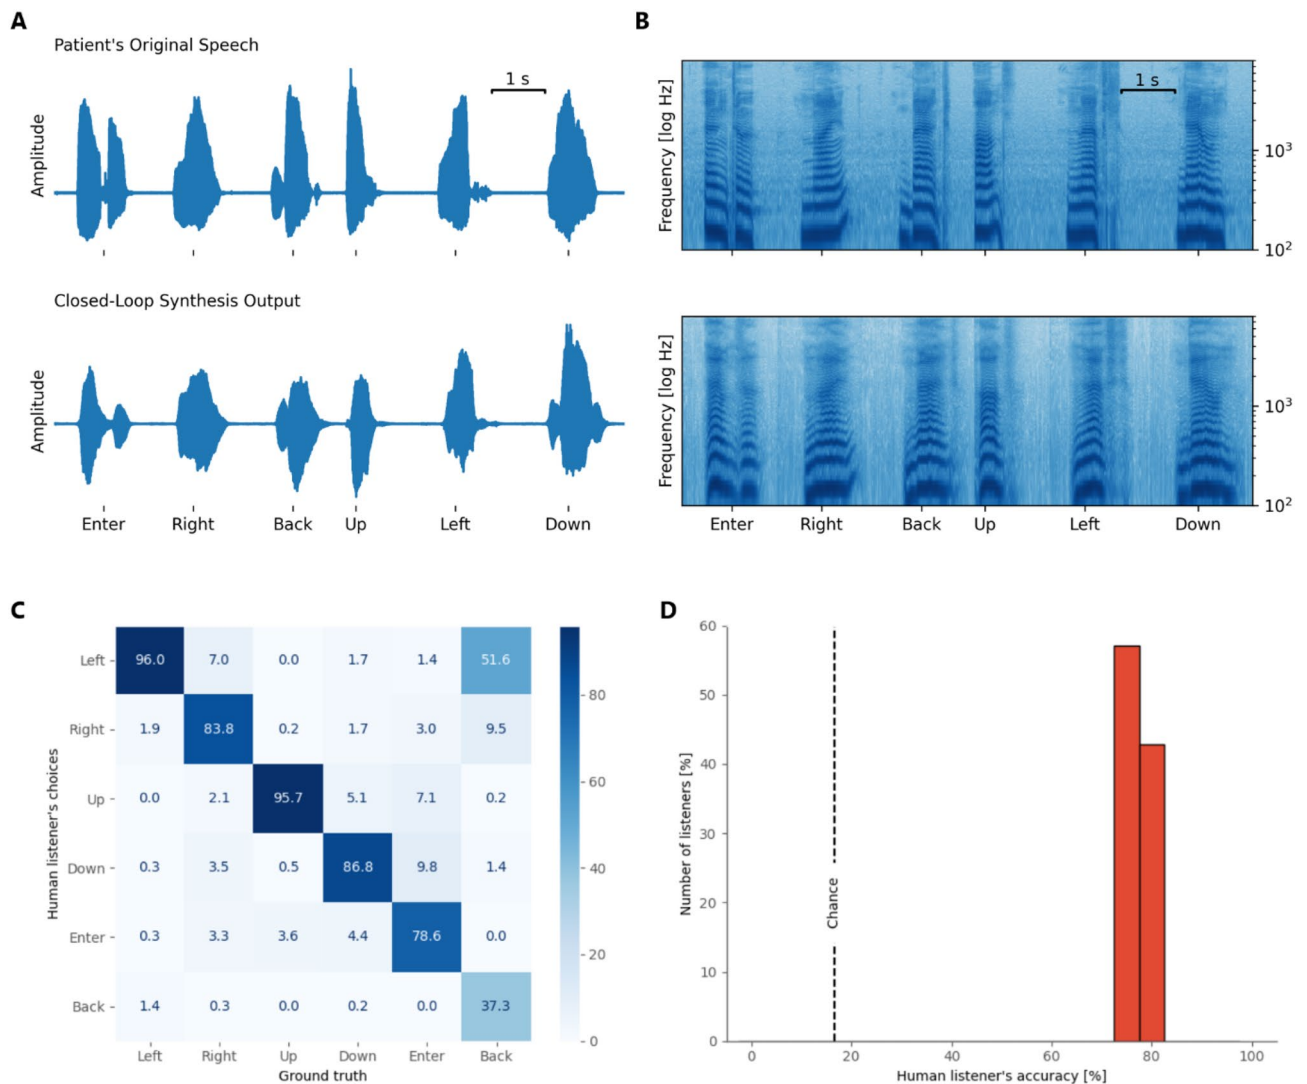


Figure 2. Evaluation of the synthesized words. (A) Visual example of time-aligned original and reconstructed acoustic speech waveforms and their spectral representations (B) for 6 words that were recorded during one of the closed-loop sessions. Speech spectrograms are shown between 100 and 8000 Hz with a logarithmic frequency range to emphasize formant frequencies. (C) The confusion matrix between human listeners and ground truth. (D) Distribution of accuracy scores from all who performed the listening test for the synthesized speech samples. Dashed line shows chance performance (16.7%).

participant (see Supplementary Video 1 for an excerpt). Other experiment parameters were not changed. All synthesized words were played back on loudspeakers while simultaneously recorded for evaluation.

To assess the intelligibility of the synthesized words, we conducted listening tests in which human listeners played back individual samples of the synthesized words and selected the word that most closely resembled each sample. Additionally, we mixed in samples that contained the originally spoken words. This allowed us to assess the quality of the participant's natural speech. We recruited a cohort of 21 native English speakers to listen to all samples that were produced during our 3 closed-loop sessions. Out of 180 samples, we excluded 2 words because the nVAD model did not detect speech activity and therefore no speech output was produced by the decoding model. We also excluded a few cases where speech activity was falsely detected by the nVAD model, which resulted in synthesized silence and remained unnoticed to the participant.

Overall, human listeners achieved an accuracy score of 80%, indicating that the majority of synthesized words could be correctly and reliably recognized. Figure 2C presents the confusion matrix regarding only the synthesized samples where the ground truth labels and human listener choices are displayed on the X- and Y-axes respectively. The confusion matrix shows that human listeners were able to recognize all but one word at very high rates. "Back" was recognized at low rates, albeit still above chance, and was most often mistaken for "Left". This could have been due in part to the close proximity of the vowel formant frequencies for these two words. The participant's weak tongue movements may have deemphasized the acoustic discriminability of these words, in turn resulting in the vocoder synthesizing a version of "back" that was often indistinct from "left". In contrast, the confusion matrix also shows that human listeners were confident in distinguishing the words "Up" and "Left". The decoder synthesized an intelligible but incorrect word in only 4% of the cases, and all listeners accurately recognized the incorrect word. Note that all keywords in the vocabulary were chosen for intuitive command and control of a computer interface, for example a communication board, and were not designed to be easily discriminable for BCI applications.

Figure 2D summarizes individual accuracy scores from all human listeners from the listening test in a histogram. All listeners recognized between 75 and 84% of the synthesized words. All human listeners achieved accuracy scores above chance (16.7%). In contrast, when tested on the participant's natural speech, our human listeners correctly recognized almost all samples of the 6 keywords (99.8%).

Anatomical and temporal contributions

In order to understand which cortical areas contributed to identification of speech segments, we conducted a saliency analysis³⁷ to reveal the underlying dynamics in high-gamma activity changes that explain the binary decisions made by our nVAD model. We utilized a method from the image processing domain³⁸ that queries spatial information indicating which pixels have contributed to a classification task. In our case, this method ranked individual high-gamma features over time by their influence on the predicted speech onsets (PSO). We defined the PSO as the first occurrence when the nVAD model identified spoken speech and neural data started to get buffered before being forwarded to the decoding model. The absolute values of their gradients allowed interpretations of which contributions had the highest or lowest impact on the class scores from anatomical and temporal perspectives.

The general idea is illustrated in Fig. 3B. In a forward pass, we first estimated for each trial the PSO by propagating through each time step until the nVAD model made a positive prediction. From here, we then applied backpropagation through time to compute all gradients with respect to the model's input high-gamma features. Relevance scores $|R|$ were computed by taking the absolute value of each partial derivative and the maximum value across time was used as the final score for each electrode³⁸. Note that we only performed backpropagation through time for each PSO, and not for whole speech segments.

Results from the saliency analysis are shown in Fig. 3A. For each channel, we display the PSO-specific relevance scores by encoding the maximum magnitude of the influence in the size of the circles (bigger circles mean stronger influence on the predictions), and the temporal occurrence of that maximum in the respective color coding (lighter electrodes have their maximal influence on the PSO earlier). The color bar at the bottom limits the temporal influence to -400 ms prior to PSO, consistent with previous reports about speech planning³⁹ and articulatory representations¹⁹. The saliency analysis showed that the nVAD model relied on a broad network of electrodes covering motor, premotor and somatosensory cortices whose collective changes in the high-gamma activity were relevant for identifying speech. Meanwhile, voice activity information encoded in the dorsal laryngeal area (highlighted electrodes in the upper grid in Fig. 3A)¹⁹ only mildly contributed to the PSO.

Figure 3C shows relevance scores over a time period of 1 s prior to PSO for 3 selected electrodes that strongly contributed to predicting speech onsets. In conjunction with the color coding from Fig. 3A, the temporal associations were consistent with previous studies that examined phoneme decoding over fixed window sizes of 400 ms¹⁸ and 500 ms^{40,41} around speech onset times, suggesting that the nVAD model benefited from neural activity during speech planning and phonological processing³⁹ when identifying speech onset. We hypothesize that the decline in the relevance scores after -200 ms can be explained by the fact that voice activity information might have already been stored in the long short-term memory of the nVAD model and thus changes in neural activity beyond this time had less influence on the prediction.

Discussion

Here we demonstrate the feasibility of a closed-loop BCI that is capable of online synthesis of intelligible words using intracranial recordings from the speech cortex of an ALS clinical trial participant. Recent studies^{10,11,13,27} suggest that deep learning techniques are a viable tool to reconstruct acoustic speech from ECoG signals. We found an approach consisting of three consecutive RNN architectures that identify and transform neural speech

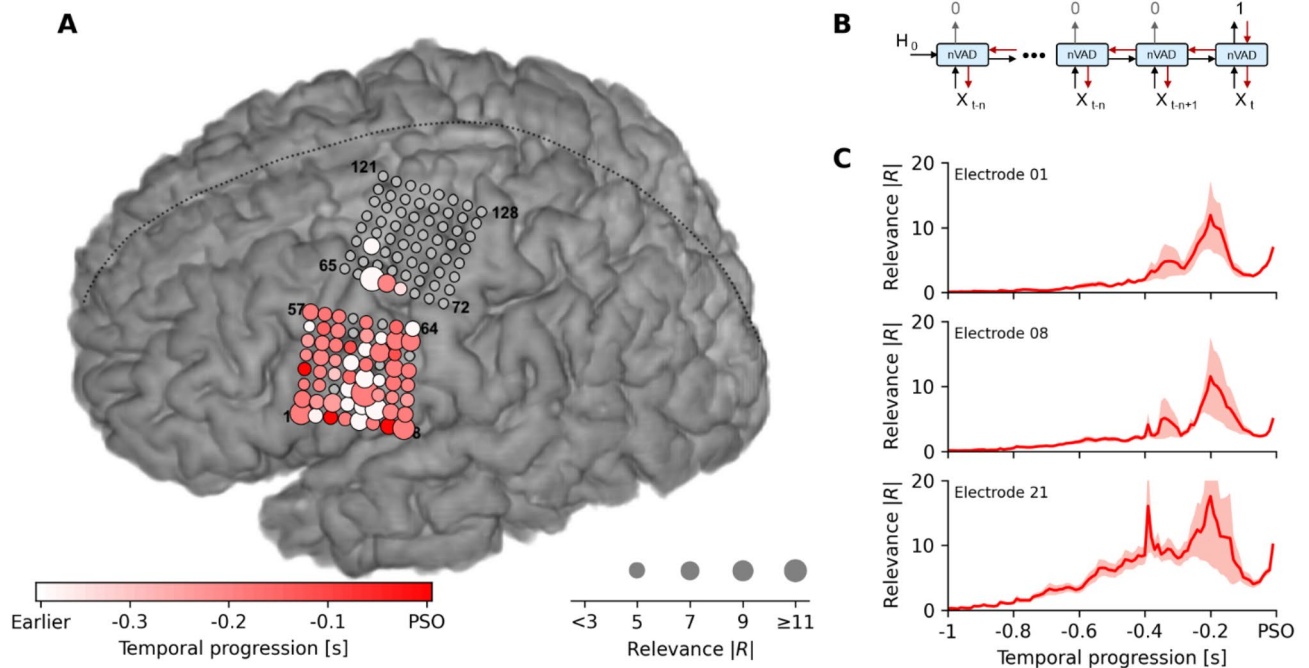


Figure 3. Changes in high-gamma activity across motor, premotor and somatosensory cortices trigger detection of speech output. **(A)** Saliency analysis shows that changes in high-gamma activity predominantly from 300 to 100 ms prior to predicted speech onset (PSO) strongly influenced the nVAD model's decision. Electrodes covering motor, premotor and somatosensory cortices show the impact of model decisions, while electrodes covering the dorsal laryngeal area only modestly added information to the prediction. Grey electrodes were either not used, bad channels or had no notable contributions. **(B)** Illustration of the general procedure on how relevance scores were computed. For each time step t , relevance scores were computed by backpropagation through time across all previous high-gamma frames X_r . Predictions of 0 correspond to no-speech, while 1 represents speech frames. **(C)** Temporal progression of mean magnitudes of the absolute relevance score in 3 selected channels that strongly contributed to PSOs. Shaded areas reflect the standard error of the mean ($N=60$). Units of the relevance scores are in 10^{-3} .

correlates into an acoustic waveform that can be streamed over the loudspeaker as neurofeedback, resulting in an 80% intelligibility score on a closed-vocabulary, keyword reading task.

The majority of human listeners were able to correctly recognize most synthesized words. All words from the closed vocabulary were chosen for a prior study²⁸ that explored speech decoding for intuitive control of a communication board rather than being constructed to elicit discriminable neural activity that benefits decoder performance. The listening tests suggest that the words “Left” and “Back” were responsible for the majority of misclassified words. These words share very similar articulatory features, and our participant's speech impairments likely made these words less discriminable in the synthesis process.

Saliency analysis showed that our nVAD approach used information encoded in the high-gamma band across predominantly motor, premotor and somatosensory cortices, while electrodes covering the dorsal laryngeal area only marginally contributed to the identification of speech onsets. In particular, neural changes previously reported to be important for speech planning and phonological processing^{19,39} appeared to have a profound impact. Here, the analysis indicates that our nVAD model learned a proper representation of spoken speech processes, providing a connection between neural patterns learned by the model and the spatio-temporal dynamics of speech production.

Our participant was chronically implanted with 128 subdural ECoG electrodes, roughly half of which covered cortical areas where similar high-gamma responses have been reliably elicited during overt speech^{18,19,40,42} and have been used for offline decoding and reconstruction of speech^{10,11}. This study and others like it^{24,27,43,44} explored the potential of ECoG-based BCIs to augment communication for individuals with motor speech impairments due to a variety of neurological disorders, including ALS and brainstem stroke. A potential advantage of ECoG for BCI is the stability of signal quality over long periods of time⁴⁵. In a previous study of an individual with locked-in syndrome due to ALS, a fully implantable ECoG BCI with fewer electrodes provided a stable switch for a spelling application over a period of more than 3 years⁴⁶. Similarly, Rao et al. reported robust responses for ECoG recordings over the speech-auditory cortex for two drug-resistant epilepsy patients over a period of 1.5 years⁴⁷. More recently, we showed that the same clinical trial participant could control a communication board with ECoG decoding of self-paced speech commands over a period of 3 months without retraining or recalibration²⁸. The speech synthesis approach we demonstrated here used training data from five and a half months prior to testing and produced similar results over 3 separate days of testing, with recalibration but no retraining in each session. These findings suggest that the correspondence between neural activity in ventral sensorimotor cortex and speech acoustics were not significantly changed over this time period. Although longitudinal testing over

longer time periods will be needed to explicitly test this, our findings provide additional support for the stability of ECoG as a BCI signal source for speech synthesis.

Our approach used a speech synthesis model trained on neural data acquired during overt speech production. This constrains our current approach to patients with speech motor impairments in which vocalization is still possible and in which speech may still be intelligible. Given the increasing use of voice banking among people living with ALS, it may also be possible to improve the intelligibility of synthetic speech using an approach similar to ours, even in participants with unintelligible or absent speech. This speech could be utilized as a surrogate but would require careful alignment to speech attempts. Likewise, the same approach could be used with a generic voice, though this would not preserve the individual's speech characteristics. Here our results were achieved without the added challenge of absent ground truth, but they serve as an important demonstration that if adequate alignment is achieved, direct synthesis of acoustic speech from ECoG is feasible, accurate, and stable, even in a person with dysarthria due to ALS. Nevertheless, it remains to be seen how long our approach will continue to produce intelligible speech as our patient's neural responses and articulatory impairments change over time due to ALS. Previous studies of long-term ECoG signal stability and BCI performance in patients with more severe motor impairments suggest that this may be possible^{3,48}.

Although our approach allowed for online, closed-loop production of synthetic speech that preserved our participant's individual voice characteristics, the bidirectional LSTM imposed a delay in the audible feedback until after the patient spoke each word. We considered this delay to be not only acceptable, but potentially desirable, given our patient's speech impairments and the likelihood of these impairments worsening in the future due to ALS. Although normal speakers use immediate acoustic feedback to tune their speech motor output⁴⁹, individuals with progressive motor speech impairments are likely to reach a point at which there is a significant, and distracting, mismatch between the subject's speech and the synthetic speech produced by the BCI. In contrast, providing acoustic feedback immediately after each utterance gives the user clear and uninterrupted output that they can use to improve subsequent speech attempts, if necessary.

While our results are promising, the approach used here did not allow for synthesis of unseen words. The bidirectional architecture of the decoding model learned variations of the neural dynamics of each word and was capable of recovering their acoustic representations from corresponding sequences of high-gamma frames. This approach did not capture more fine-grained and isolated part-of-speech units, such as syllables or phonemes. However, previous research^{11,27} has shown that speech synthesis approaches based on bidirectional architectures can generalize to unseen elements that were not part of the training set. Future research will be needed to expand the limited vocabulary used here, and to explore to what extent similar or different approaches are able to extrapolate to words that are not in the vocabulary of the training set.

Our demonstration here builds on previous seminal studies of the cortical representations for articulation and phonation^{19,32,40} in epilepsy patients implanted with similar subdural ECoG arrays for less than 30 days. These studies and others using intraoperative recordings have also supported the feasibility of producing synthetic speech from ECoG high-gamma responses^{10,11,33}, but these demonstrations were based on offline analysis of ECoG signals that were previously recorded in subjects with normal speech, with the exception of the work by Metzger et al.²⁷ Here, a participant with impaired articulation and phonation was able to use a chronically implanted investigational device to produce acoustic speech that retained his unique voice characteristics. This was made possible through online decoding of ECoG high-gamma responses, using an algorithm trained on data collected months before. Notwithstanding the current limitations of our approach, our findings here provide a promising proof-of-concept that ECoG BCIs utilizing online speech synthesis can serve as alternative and augmentative communication devices for people living with ALS. Moreover, our findings should motivate continued research on the feasibility of using BCIs to preserve or restore vocal communication in clinical populations where this is needed.

Materials and methods

Participant

Our participant was a male native English speaker in his 60s with ALS who was enrolled in a clinical trial (NCT03567213), approved by the Johns Hopkins University Institutional Review Board (IRB) and by the FDA (under an investigational device exemption) to test the safety and preliminary efficacy of a brain-computer interface composed of subdural electrodes and a percutaneous connection to external EEG amplifiers and computers. All experiments conducted in this study complied with all relevant guidelines and regulations, and were performed according to a clinical trial protocol approved by the Johns Hopkins IRB. Diagnosed with ALS 8 years prior to implantation, our participant's motor impairments had chiefly affected bulbar and upper extremity muscles and had resulted in motor impairments sufficient to render continuous speech mostly unintelligible (though individual words were intelligible), and to require assistance with most activities of daily living. Our participant's ability to carry out activities of daily living were assessed using the ALSFRS-R measure⁵⁰, resulting in a score of 26 out of 48 possible points (speech was rated at 1 point, see Supplementary Data S5). Furthermore, speech intelligibility and speaking rate were evaluated by a certified speech-language pathologist, whose detailed assessment may be found in the Supplementary Note. The participant gave informed consent after being counseled about the nature of the research and implant-related risks and was implanted with the study device in July 2022. Additionally, the participant gave informed consent for use of his audio and video recordings in publications of the study results.

Study device and implantation

The study device was composed of two 8 × 8 subdural electrode grids (PMT Corporation, Chanhassen, MN) connected to a percutaneous 128-channel Neuroport pedestal (Blackrock Neurotech, Salt Lake City, UT). Both

subdural grids contained platinum-iridium disc electrodes (0.76 mm thickness, 2-mm diameter exposed surface) with 4 mm center-to-center spacing and a total surface area of 12.11 cm² (36.6 mm × 33.1 mm).

The study device was surgically implanted during a standard awake craniotomy with a combination of local anesthesia and light sedation, without neuromuscular blockade. The device's ECoG grids were placed on the pial surface of sensorimotor representations for speech and upper extremity movements in the left hemisphere. Careful attention was made to assure that the scalp flap incision was well away from the external pedestal. Cortical representations were targeted using anatomical landmarks from pre-operative structural (MRI) and functional imaging (fMRI), in addition to somatosensory evoked potentials measured intraoperatively. Two reference wires attached to the Neuroport pedestal were implanted in the subdural space on the outward facing surface of the subdural grids. The participant was awoken during the craniotomy to confirm proper functioning of the study device and final placement of the two subdural grids. For this purpose, the participant was asked to repeatedly speak a single word as event-related ECoG spectral responses were noted to verify optimal placement for the implanted electrodes. On the same day, the participant had a post-operative CT which was then co-registered to a pre-operative MRI to verify the anatomical locations of the two grids.

Data recording

During all training and testing sessions, the Neuroport pedestal was connected to a 128-channel NeuroPlex-E headstage that was in turn connected by a mini-HDMI cable to a NeuroPort Biopotential Signal Processor (Blackrock Neurotech, Salt Lake City, UT, USA) and external computers. We acquired neural signals at a sampling rate of 1000 Hz.

Acoustic speech was recorded through an external microphone (BETA[®] 58A, SHURE, Niles, IL) in a room isolated from external acoustic and electronic noise, then amplified and digitized by an external audio interface (H6-audio-recorder, Zoom Corporation, Tokyo, Japan). The acoustic speech signal was split and forwarded to: (1) an analog input of the NeuroPort Biopotential Signal Processor (NSP) to be recorded at the same frequency and in synchrony with the neural signals, and (2) the testing computer to capture high-quality (48 kHz) recordings. We applied cross-correlation to align the high-quality recordings with the synchronized audio signal from the NSP.

Experiment recordings and task design

Each recording day began with a syllable repetition task to acquire cortical activity to be used for baseline normalization. Each syllable was audibly presented through a loudspeaker, and the participant was instructed to recite the heard stimulus by repeating it aloud. Stimulus presentation lasted for 1 s, and trial duration was set randomly in the range of 2.5 s and 3.5 s with a step size of 80 ms. In the syllable repetition task, the participant was instructed to repeat 12 consonant–vowel syllables (Supplementary Table S4), in which each syllable was repeated 5 times. We extracted high-gamma frames from all trials to compute for each day the mean and standard deviation statistics for channel-specific normalization.

To collect data for training our nVAD and speech decoding model, we recorded ECoG during multiple blocks of a speech production task over a period of 6 weeks. During the task, the participant read aloud single words that were prompted on a computer screen, interrupted occasionally by a silence trial in which the participant was instructed to say nothing. The words came from a closed vocabulary of 6 words ("Left", "Right", "Up", "Down", "Enter", "Back", and "... for silence) that were chosen for a separate study in which these spoken words were decoded from ECoG to control a communication board²⁸. In each block, there were ten repetitions of each word (60 words in total) that appeared in a pseudo-randomized order by having a fixed set of seeds to control randomization orders. Each word was shown for 2 s per trial with an intertrial interval of 3 s. The participant was instructed to read the prompted word aloud as soon as it appeared. Because his speech was slow, effortful, and dysarthric, the participant may have sometimes used some of the intertrial interval to complete word production. However, offline analysis verified at least 1 s between the end of each spoken word and the beginning of the next trial, assuring that enough time had passed to avoid ECoG high-gamma responses leaking into subsequent trials. In each block, neural signals and audibly vocalized speech were acquired in parallel and stored to disc using BCI2000³¹.

We recorded training, validation, and test data for 10 days, and deployed our approach for synthesizing speech online five and a half months later. During the online task, the synthesized output was played to the participant while he performed the same keyword reading task as in the training sessions. The feedback from each synthesized word began after he spoke the same word, avoiding any interference with production from the acoustic feedback. The validation dataset was used for finding appropriate hyperparameters to train both nVAD and the decoding model. The test set was used to validate final model generalizability before online sessions. We also used the test set for the saliency analysis. In total, the training set was comprised of 1570 trials that aggregated to approximately 80 min of data (21.8 min are pure speech), while the validation and test set contained 70 trials each with around 3 min of data (0.9 min pure speech). The data in each of these datasets were collected on different days, so that no baseline or other statistics in the training set leaked into the validation or test set.

Signal processing and feature extraction

Neural signals were transformed into broadband high-gamma power features that have been previously reported to closely track the timing and location of cortical activation during speech and language processes^{42,52}. In this feature extraction process, we first re-referenced all channels within each 64-contact grid to a common-average reference (CAR filtering), excluding channels with poor signal quality in any training session. Next, we selected all channels that had previously shown significant high-gamma responses during the syllable repetition task described above. This included 64 channels (Supplementary Fig. S2, channels with blue outlines) across motor,

premotor and somatosensory cortices, including the dorsal laryngeal area. From here, we applied two IIR Butterworth filters (both with filter order 8) to extract the high-gamma band in the range of 70 to 170 Hz while subsequently attenuating the first harmonic (118–122 Hz) of the line noise. For each channel, we computed logarithmic power features based on windows with a fixed length of 50 ms and a frameshift of 10 ms. To estimate speech-related increases in broadband high-gamma power, we normalized each feature by the day-specific statistics of the high-gamma power features accumulated from the syllable repetition task.

For the acoustic recordings of the participant's speech, we downsampled the time-aligned high-quality microphone recordings from 48 to 16 kHz. From here, we padded the acoustic data by 16 ms to account for the shift introduced by the two filters on the neural data and estimated the boundaries of speech segments using an energy-based voice activity detection algorithm⁵³. Likewise, we computed acoustic features in the LPC coefficient space through the encoding functionality of the LPCNet vocoder. Both voice activity detection and LPC feature encoding were configured to operate on 10 ms frameshifts to match the number of samples from the broadband high-gamma feature extraction pipeline.

Network architectures

Our proposed approach relied on three recurrent neural network architectures: (1) a unidirectional model that identified speech segments from the neural data, (2) a bidirectional model that translated sequences of speech-related high-gamma activity into corresponding sequences of LPC coefficients representing acoustic information, and (3) LPCNet³⁶, which converted those LPC coefficients into an acoustic speech signal.

The network architecture of the unidirectional nVAD model was inspired by Zen et al.⁵⁴ in using a stack of two LSTM layers with 150 units each, followed by a linear fully connected output layer with two units representing speech or non-speech class target logits (Fig. 4). We trained the unidirectional nVAD model using truncated backpropagation through time (BPTT)⁵⁵ to keep the costs of single parameter updates manageable. We initialized this algorithm's hyperparameters k_1 and k_2 to 50 and 100 frames of high-gamma activity, respectively, such that the unfolding procedure of the backpropagation step was limited to 100 frames (1 s) and repeated every 50 frames (500 ms). Dropout was used as a regularization method with a probability of 50% to counter overfitting effects⁵⁶. Comparison between predicted and target labels was determined by the cross-entropy loss. We limited the network training using an early stopping mechanism that evaluated after each epoch the network performance on a held-out validation set and kept track of the best model weights by storing the model weights only when the frame-wise accuracy score was bigger than before. The learning rate of the stochastic gradient descent optimizer was dynamically adjusted in accordance with the RMSprop formula⁵⁷ with an initial learning rate of 0.001. Using this procedure, the unidirectional nVAD model was trained for 27,975 update steps, achieving a frame-wise accuracy of 93.4% on held-out validation data. The architecture of the nVAD model had 311,102 trainable weights.

The network architecture of the bidirectional decoding model had a very similar configuration to the unidirectional nVAD but employed a stack of bidirectional LSTM layers for sequence modelling¹¹ to include past and future contexts. Since the acoustic space of the LPC components was continuous, we used a linear fully connected output layer for this regression task. Figure 4 contains an illustration of the network architecture of the decoding model. In contrast to the unidirectional nVAD model, we used standard BPTT to account for both past and future contexts within each extracted segment identified as spoken speech. The architecture of the decoding model had 378,420 trainable weights and was trained for 14,130 update steps using a stochastic gradient descent optimizer. The initial learning rate was set to 0.001 and dynamically updated in accordance with the RMSProp

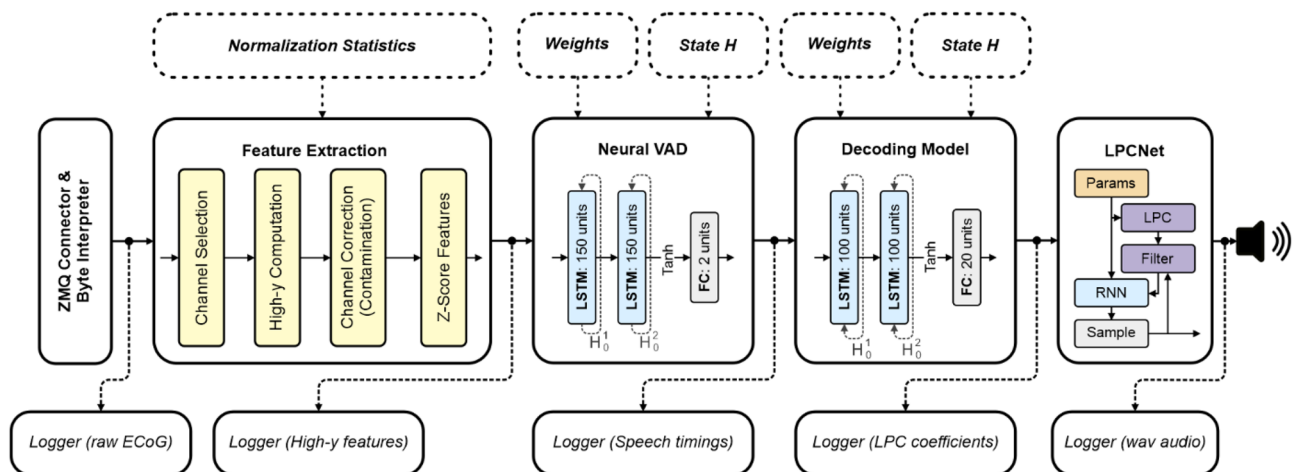


Figure 4. System overview of the closed-loop architecture. The computational graph is designed as a directed acyclic network. Solid shapes represent ezmsg units, dotted ones represent initialization parameters. Each unit is responsible for a self-contained task and distributes their output to all its subscribers. Logger units run in separate processes to not interrupt the main processing chain for synthesizing speech.

formula. Again, we used dropout with a 50% probability and employed an early stopping mechanism that only updated model weights when the loss on the held-out validation set was lower than before.

Both the unidirectional nVAD and the bidirectional decoding model were implemented within the PyTorch framework. For LPCNet, we used the C-implementation and pretrained model weights by the original authors and communicated with the library via wrapper functions through the Cython programming language.

Closed-loop architecture

Our closed-loop architecture was built upon *ezmsg*, a general-purpose framework which enables the implementation of streaming systems in the form a directed acyclic network of connected units, which communicate with each other through a publish/subscribe software engineering pattern using asynchronous coroutines. Here, each unit represents a self-contained operation which receives many inputs, and optionally propagates its output to all its subscribers. A unit consists of a settings and state class for enabling initial and updatable configurations and has multiple input and output connection streams to communicate with other nodes in the network. Figure 4 shows a schematic overview of the closed-loop architecture. ECoG signals were received by connecting to BCI2000 via a custom ZeroMQ (ZMQ) networking interface that sent packages of 40 ms over the TCP/IP protocol. From here, each unit interacted with other units through an asynchronous message system that was implemented on top of a shared-memory publish-subscribe multi-processing pattern. Figure 4 shows that the closed-loop architecture was comprised of 5 units for the synthesis pipeline, while employing several additional units that acted as loggers and wrote intermediate data to disc.

In order to play back the synthesized speech during closed-loop sessions, we wrote the bytes of the raw PCM waveform to standard output (stdout) and reinterpreted them by piping them into SoX. We implemented our closed-loop architecture in Python 3.10. To keep the computational complexity manageable for this streamlined application, we implemented several functionalities, such as ringbuffers or specific calculations in the high-gamma feature extraction, in Cython.

Contamination analysis

Overt speech production can cause acoustic artifacts in electrophysiological recordings, allowing learning machines such as neural networks to rely on information that is likely to fail once deployed—a phenomenon widely known as Clever Hans⁵⁸. We used the method proposed by Roussel et al.⁵⁹ to assess the risk that our ECoG recordings had been contaminated. This method compares correlations between neural and acoustic spectrograms to determine a contamination index which describes the average correlation of matching frequencies. This contamination index is compared to the distribution of contamination indices resulting from randomly permuting the rows and columns of the contamination matrix—allowing statistical analysis of the risk when assuming that no acoustic contamination is present.

For each recording day among the train, test and validation set, we analyzed acoustic contamination in the high-gamma frequency range. We identified 1 channel (Channel 46) in our recordings that was likely contaminated during 3 recording days (D_5 , D_6 , and D_7), and we corrected this channel by taking the average of high-gamma power features from neighboring channels (8-neighbour configuration, excluding the bad channel 38). A detailed report can be found in Supplementary Fig. S1, where each histogram corresponds to the distribution of permuted contamination matrices, and colored vertical bars indicate the actual contamination index, where green and red indicate the statistical criterion threshold (green: $p > 0.05$, red: $p \leq 0.05$). After excluding the neural data from channel 46, Roussel's method suggested that the null hypothesis could be rejected, and thus we concluded that no acoustic speech has interfered with neural recording.

Listening test

We conducted a forced-choice listening test similar to Herff et al.¹⁴ in which 21 native English speakers evaluated the intelligibility of the synthesized output and the originally spoken words. Listeners were asked to listen to one word at a time and select which word out of the six options most closely resembled it. Here, the listeners had the opportunity to listen to each sample many times before submitting a choice. We implemented the listening test on top of the BeagleJS framework⁶⁰. All words that were either spoken or synthesized during the 3 closed-loop sessions were included in the listening test, but were randomly sampled from a uniform distribution for unique randomized sequences across listeners. Supplementary Fig. S3 provides a screenshot of the interface with which the listeners were working.

All human listeners were only recruited through indirect means such as IRB-approved flyers placed on campus sites and had no direct connection to the PI. Anonymous demographic data was collected at the end of the listening test asking for age and preferred gender. Overall, recruited participants were 23.8% male and 61.9% female (14% other or preferred not to answer) ranging between 18 to 30 years old.

Statistical analysis

Original and reconstructed speech spectrograms were compared using Pearson's correlation coefficients for 80 mel-scaled spectral bins. For this, we transformed original and reconstructed waveforms into the spectral domain using the short-time Fourier transform (window size: 50 ms, frameshift: 10 ms, window function: Hanning), applied 80 triangular filters to focus only on perceptual differences for human listeners⁶¹, and Gaussianized the distribution of the acoustic space using the natural logarithm. Pearson correlation scores were calculated for each sample by averaging the correlation coefficients across frequency bins. The 95% confidence interval (two-sided) was used in the feature selection procedure while the z-criterion was Bonferroni corrected across time points. Lower and upper bounds for all channels and time points can be found in the supplementary data. Contamination analysis is based on permutation tests that use t-tests as their statistical criterion with a Bonferroni

corrected significance level of $\alpha = 0.05/N$, where N represents the number of frequency bins multiplied by the number of selected channels.

Overall, we used the SciPy stats package (version 1.10.1) for statistical evaluation, but the contamination analysis has been done in Matlab with the statistics and machine learning toolbox (version 12.4).

Data availability

Neural data and anonymized speech audio are publicly available at <http://www.osf.io/49rt7/>. This includes experiment recordings used as training data and experiment runs from our closed-loop sessions. Additionally, we also included supporting data used for rendering the figures in the main text and in the supplementary material.

Code availability

Corresponding source code for the closed-loop BCI and scripts for generating figures can be obtained from the official Crone Lab Github page at: <https://github.com/cronelab/delayed-speech-synthesis>. This includes source files for training, inference, and data analysis/evaluation. The ezmsg framework can be obtained from <https://github.com/iscoe/ezmsg>.

Received: 19 October 2023; Accepted: 21 April 2024

Published online: 26 April 2024

References

- Bauer, G., Gerstenbrand, F. & Rimpl, E. Varieties of the locked-in syndrome. *J. Neurol.* **221**, 77–91 (1979).
- Smith, E. & Delargy, M. Locked-in syndrome. *BMJ* **330**, 406–409 (2005).
- 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).
- Chaudhary, U. *et al.* Spelling interface using intracortical signals in a completely locked-in patient enabled via auditory neurofeedback training. *Nat. Commun.* **13**, 1236 (2022).
- Pandarinath, C. *et al.* High performance communication by people with paralysis using an intracortical brain–computer interface. *eLife* **6**, e18554 (2017).
- Willett, F. R., Avansino, D. T., Hochberg, L. R., Henderson, J. M. & Shenoy, K. V. High-performance brain-to-text communication via handwriting. *Nature* **593**, 249–254 (2021).
- Oxley, T. J. *et al.* Motor neuroprosthesis implanted with neurointerventional surgery improves capacity for activities of daily living tasks in severe paralysis: First in-human experience. *J. NeuroInterventional Surg.* **13**, 102–108 (2021).
- Chang, E. F. & Anumanchipalli, G. K. Toward a speech neuroprosthesis. *JAMA* **323**, 413–414 (2020).
- Herff, C. *et al.* Towards direct speech synthesis from ECoG: A pilot study. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* 1540–1543 (2016).
- Angrick, M. *et al.* Speech synthesis from ECoG using densely connected 3D convolutional neural networks. *J. Neural Eng.* **16**, 036019 (2019).
- Anumanchipalli, G. K., Chartier, J. & Chang, E. F. Speech synthesis from neural decoding of spoken sentences. *Nature* **568**, 493–498 (2019).
- Wairagkar, M., Hochberg, L. R., Brandman, D. M. & Stavisky, S. D. Synthesizing speech by decoding intracortical neural activity from dorsal motor cortex. In *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)* 1–4 (2023).
- Kohler, J. *et al.* Synthesizing speech from intracranial depth electrodes using an encoder-decoder framework. *Neurons Behav. Data Anal. Theory* <https://doi.org/10.51628/001c.57524> (2022).
- Herff, C. *et al.* Generating natural, intelligible speech from brain activity in motor, premotor, and inferior frontal cortices. *Front. Neurosci.* <https://doi.org/10.3389/fnins.2019.01267> (2019).
- Wilson, G. H. *et al.* Decoding spoken English from intracortical electrode arrays in dorsal precentral gyrus. *J. Neural Eng.* **17**, 066007 (2020).
- Kanas, V. G. *et al.* Joint spatial-spectral feature space clustering for speech activity detection from ECoG signals. *IEEE Trans. Biomed. Eng.* **61**, 1241–1250 (2014).
- Soroush, P. Z., Angrick, M., Shih, J., Schultz, T. & Krusienski, D. J. Speech activity detection from stereotactic EEG. In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* 3402–3407 (2021).
- 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).
- Bouchard, K. E., Mesgarani, N., Johnson, K. & Chang, E. F. Functional organization of human sensorimotor cortex for speech articulation. *Nature* **495**, 327–332 (2013).
- Bouchard, K. E. & Chang, E. F. Neural decoding of spoken vowels from human sensory-motor cortex with high-density electrocorticography. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* 6782–6785 (2014).
- Kellis, S. *et al.* Decoding spoken words using local field potentials recorded from the cortical surface. *J. Neural Eng.* **7**, 056007 (2010).
- Mugler, E. M., Goldrick, M., Rosenow, J. M., Tate, M. C. & Slutzky, M. W. Decoding of articulatory gestures during word production using speech motor and premotor cortical activity. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* 5339–5342 (2015).
- Mugler, E. M. *et al.* Differential representation of articulatory gestures and phonemes in precentral and inferior frontal gyri. *J. Neurosci.* **38**, 9803–9813 (2018).
- Moses, D. A. *et al.* Neuroprosthesis for decoding speech in a paralyzed person with anarthria. *N. Engl. J. Med.* **385**, 217–227 (2021).
- Willett, F. R. *et al.* A high-performance speech neuroprosthesis. *Nature* **620**, 1031–1036 (2023).
- Guenther, F. H. *et al.* A wireless brain–machine interface for real-time speech synthesis. *PLoS ONE* **4**, e8218 (2009).
- Metzger, S. L. *et al.* A high-performance neuroprosthesis for speech decoding and avatar control. *Nature* **620**, 1037–1046 (2023).
- Luo, S. *et al.* Stable decoding from a speech BCI enables control for an individual with ALS without recalibration for 3 months. *Adv. Sci.* **10**, 2304853 (2023).
- Cooney, C., Folli, R. & Coyle, D. Neurolinguistics research advancing development of a direct-speech brain–computer interface. *iScience* **8**, 103–125 (2018).
- Herff, C. & Schultz, T. Automatic speech recognition from neural signals: A focused review. *Front. Neurosci.* <https://doi.org/10.3389/fnins.2016.00429> (2016).
- Dash, D. *et al.* Neural Speech Decoding for Amyotrophic Lateral Sclerosis, 2782–2786 (2020). <https://doi.org/10.21437/Interspeech.2020-3071>.

32. Chartier, J., Anumanchipalli, G. K., Johnson, K. & Chang, E. F. Encoding of articulatory kinematic trajectories in human speech sensorimotor cortex. *Neuron* **98**, 1042–1054.e4 (2018).
33. Akbari, H., Khalighinejad, B., Herrero, J. L., Mehta, A. D. & Mesgarani, N. Towards reconstructing intelligible speech from the human auditory cortex. *Sci. Rep.* **9**, 874 (2019).
34. Moore, B. An introduction to the psychology of hearing: Sixth edition. In *An Introduction to the Psychology of Hearing* (Brill, 2013).
35. Taylor, P. *Text-to-Speech Synthesis* (Cambridge University Press, 2009).
36. Valin, J.-M. & Skoglund, J. LPCNET: Improving neural speech synthesis through linear prediction. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 5891–5895 (2019).
37. Montavon, G., Samek, W. & Müller, K.-R. Methods for interpreting and understanding deep neural networks. *Digit. Signal Process.* **73**, 1–15 (2018).
38. Simonyan, K., Vedaldi, A. & Zisserman, A. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *International Conference on Learning Representations (ICLR)* (2014).
39. Indefrey, P. the spatial and temporal signatures of word production components: A critical update. *Front. Psychol.* <https://doi.org/10.3389/fpsyg.2011.00255> (2011).
40. Ramsey, N. F. *et al.* Decoding spoken phonemes from sensorimotor cortex with high-density ECoG grids. *NeuroImage* **180**, 301–311 (2018).
41. Jiang, W., Pailla, T., Dichter, B., Chang, E. F. & Gilja, V. Decoding speech using the timing of neural signal modulation. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* 1532–1535 (2016).
42. Crone, N. E. *et al.* Electrographic gamma activity during word production in spoken and sign language. *Neurology* **57**, 2045–2053 (2001).
43. Moses, D. A., Leonard, M. K., Makin, J. G. & Chang, E. F. Real-time decoding of question-and-answer speech dialogue using human cortical activity. *Nat. Commun.* **10**, 3096 (2019).
44. Herff, C. *et al.* Brain-to-text: Decoding spoken phrases from phone representations in the brain. *Front. Neurosci.* <https://doi.org/10.3389/fnins.2015.00217> (2015).
45. Morrell, M. J. Responsive cortical stimulation for the treatment of medically intractable partial epilepsy. *Neurology* **77**, 1295–1304 (2011).
46. Pels, E. G. M. *et al.* Stability of a chronic implanted brain–computer interface in late-stage amyotrophic lateral sclerosis. *Clin. Neurophysiol.* **130**, 1798–1803 (2019).
47. Rao, V. R. *et al.* Chronic ambulatory electrocorticography from human speech cortex. *NeuroImage* **153**, 273–282 (2017).
48. Silversmith, D. B. *et al.* Plug-and-play control of a brain–computer interface through neural map stabilization. *Nat. Biotechnol.* **39**, 326–335 (2021).
49. Denes, P. B. & Pinson, E. *The Speech Chain* (Macmillan, 1993).
50. Cedarbaum, J. M. *et al.* The ALSFRS-R: A revised ALS functional rating scale that incorporates assessments of respiratory function. *J. Neurol. Sci.* **169**, 13–21 (1999).
51. Schalk, G., McFarland, D. J., Hinterberger, T., Birbaumer, N. & Wolpaw, J. R. BCI2000: A general-purpose brain-computer interface (BCI) system. *IEEE Trans. Biomed. Eng.* **51**, 1034–1043 (2004).
52. Leuthardt, E. *et al.* Temporal evolution of gamma activity in human cortex during an overt and covert word repetition task. *Front. Hum. Neurosci.* <https://doi.org/10.3389/fnhum.2012.00099> (2012).
53. Povey, D. *et al.* The kald speech recognition toolkit. In *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding* (IEEE Signal Processing Society, 2011).
54. Zen, H. & Sak, H. Unidirectional long short-term memory recurrent neural network with recurrent output layer for low-latency speech synthesis. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* 4470–4474 (2015).
55. Sutskever, I. *Training Recurrent Neural Networks* (University of Toronto, 2013).
56. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
57. Ruder, S. An overview of gradient descent optimization algorithms. Preprint at <https://arxiv.org/abs/1609.04747> (2016).
58. Lapuschkin, S. *et al.* Unmasking Clever Hans predictors and assessing what machines really learn. *Nat. Commun.* **10**, 1096 (2019).
59. Roussel, P. *et al.* Observation and assessment of acoustic contamination of electrophysiological brain signals during speech production and sound perception. *J. Neural Eng.* **17**, 056028 (2020).
60. Kraft, S. & Zölzer, U. BeaqleJS: HTML5 and JavaScript based framework for the subjective evaluation of audio quality. In *Linux Audio Conference* (2014).
61. Stevens, S. S., Volkman, J. & Newman, E. B. A scale for the measurement of the psychological magnitude pitch. *J. Acoust. Soc. Am.* **8**, 185–190 (1937).

Acknowledgements

Research reported in this publication was supported by the National Institute Of Neurological Disorders And Stroke of the National Institutes of Health under Award Number UH3NS114439 (PI N.E.C., co-PI N.F.R.). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Author contributions

M.A. and N.C. wrote the manuscript. M.A., S.L., Q.R. and D.C. analyzed the data. M.A. and S.S. conducted the listening test. S.L. collected the data. M.A. and G.M. implemented the code for the online decoder and the underlying framework. M.A. made the visualizations. W.A., C.G. and K.R., L.C. and N.M. conducted the surgery/medical procedure. D.T. made the speech and language assessment. F.T. handled the regulatory aspects. H.H. supervised the speech processing methodology. M.F. N.R. and N.C. supervised the study and the conceptualization. All authors reviewed and revised the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-60277-2>.

Correspondence and requests for materials should be addressed to M.A. or N.E.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024