

PAPER • OPEN ACCESS

Direct speech reconstruction from sensorimotor brain activity with optimized deep learning models

To cite this article: Julia Berezutskaya *et al* 2023 *J. Neural Eng.* **20** 056010

View the [article online](#) for updates and enhancements.

You may also like

- [Considerations for implanting speech brain computer interfaces based on functional magnetic resonance imaging](#)
F Guerreiro Fernandes, M Raemaekers, Z Freudenburg et al.
- [Implications of shared motor and perceptual activations on the sensorimotor cortex for neuroprosthetic decoding](#)
Alexander B Silva, Jessie R Liu, Vanessa R Anderson et al.
- [Word classification across speech modes from low-density electrocorticography signals](#)
Aurélien de Borman, Bob Van Dyck, Kato Van Rooy et al.



PAPER

OPEN ACCESS

RECEIVED
4 December 2022REVISED
12 July 2023ACCEPTED FOR PUBLICATION
19 July 2023PUBLISHED
20 September 2023

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Direct speech reconstruction from sensorimotor brain activity with optimized deep learning models

Julia Berezutskaya^{1,2,*} , Zachary V Freudenburg^{1,3} , Mariska J Vansteensel^{1,3} , Erik J Aarnoutse¹ , Nick F Ramsey¹ and Marcel A J van Gerven²

¹ Brain Center, Department of Neurology and Neurosurgery, University Medical Center Utrecht, Utrecht 3584 CX, The Netherlands

² Donders Center for Brain, Cognition and Behaviour, Nijmegen 6525 GD, The Netherlands

³ These authors contributed equally to this work

* Authors to whom any correspondence should be addressed.

E-mail: y.berezutskaya@umcutrecht.nl

Keywords: brain, speech, deep neural networks, brain–computer interfaces, electrocorticography, audio reconstruction, neural decoding

Supplementary material for this article is available [online](#)

Abstract

Objective. Development of brain–computer interface (BCI) technology is key for enabling communication in individuals who have lost the faculty of speech due to severe motor paralysis. A BCI control strategy that is gaining attention employs speech decoding from neural data. Recent studies have shown that a combination of direct neural recordings and advanced computational models can provide promising results. Understanding which decoding strategies deliver best and directly applicable results is crucial for advancing the field. **Approach.** In this paper, we optimized and validated a decoding approach based on speech reconstruction directly from high-density electrocorticography recordings from sensorimotor cortex during a speech production task. **Main results.** We show that (1) dedicated machine learning optimization of reconstruction models is key for achieving the best reconstruction performance; (2) individual word decoding in reconstructed speech achieves 92%–100% accuracy (chance level is 8%); (3) direct reconstruction from sensorimotor brain activity produces intelligible speech. **Significance.** These results underline the need for model optimization in achieving best speech decoding results and highlight the potential that reconstruction-based speech decoding from sensorimotor cortex can offer for development of next-generation BCI technology for communication.

1. Introduction

Due to a motor neuron disease or a brainstem stroke, people can lose all voluntary control over their muscles, including the ability to speak or use other movements for communication. Brain–computer interface (BCI) technology aims to provide a means of communication to these individuals. Recent advances in BCI research have demonstrated the potential of current approaches to record, process and analyse neural activity in order to decode speech [1–28]. This work is primarily conducted in non-disabled human subjects who participate in speech production experiments while their brain activity is recorded, with the goal to infer spoken speech from the acquired brain signals. This setup is used as a

testbed for development and validation of speech decoding frameworks prior to their use in real-world BCI applications with paralyzed individuals [29–35]. See [36–46] for reviews on BCIs for communication and speech.

Previous work demonstrates that speech decoding from brain activity is a challenging task. In order to realize a real-world speech-based BCI application, several key components of such a system need to be identified. This includes decisions on the neural recording modality (acquisition technique), the cortical areas that are most informative for decoding (neural sources), the target speech features to decode (decoding targets) and the decoding model itself. Regarding the acquisition technique, the most promising results in the field so

far appear to have come from the use of invasive recording modalities, mainly electrocorticography (ECoG) [10–12, 15, 18, 20, 23, 24]. This is because ECoG grids, and particularly those with a high spatial density (high-density ECoG; HD ECoG), provide excellent signal quality and high spatial and temporal resolution, which are key for achieving accurate decoding results. Preparatory BCI research and development can be performed in non-disabled patients with medication-resistant epilepsy who undergo temporarily implantation with ECoG electrodes (typically, for 7–10 days) for clinical monitoring of their condition and potential subsequent removal of neural sources of epilepsy.


Regarding the neural sources, among the areas most informative for decoding is the motor cortex, which coordinates voluntary muscle control. Studies in amputees and severely paralyzed individuals have shown that motor cortex activity during attempted movement of the missing limb or attempted movement by the paralyzed individual is similar to the activity during actual movement in non-disabled individuals [31, 47]. Motor-based BCI research has led to the development of various BCI applications, such as the control of a robotic arm in tetraplegia [48], a BCI speller based on attempted hand movement in a locked-in individual [31] and a BCI for speech detection and decoding in an individual with anarthria [33].

Choices of decoding targets and decoding model, however, are more difficult and varied. In principle, use of motor cortex as the neural source for decoding means that one of the most straightforward targets of decoding would be information about the movement of facial muscles involved in speech production. This is typically referred to as the kinematic articulator traces [18, 49]. However, for development of such a speech decoding approach, articulator movements need to be recorded during speech simultaneously with the brain activity, which is rarely feasible in ECoG research due to the additional burden it would inflict on the study participants. Acoustic-to-articulatory inversion models can potentially be used to facilitate decoding and sound reconstruction using kinematic traces [18, 50]. Otherwise, another system will be required to map the decoded kinematic movements to language information, such as phonemes, syllables or words [51]. One could also consider using these language labels as direct targets of decoding. This has been done with certain success especially when language models are used to correct decoding mistakes in continuous speech [33, 35]. However, accurate decoding of isolated word and phonemes from brain activity has not been trivial, especially when decoding a large number of classes [10, 30, 35]. Moreover, some ethical concerns have been raised in relation to the use of language models for BCIs [52]. An alternative candidate for the target of speech decoding is acoustic speech features

from simultaneously acquired microphone recordings of produced speech, available at the BCI development stage that relies on experiments in non-disabled participants [20, 21, 26, 27]. Such acoustic properties are fast-changing, highly repeatable speech features that have been previously shown to explain patterns of motor and premotor brain activity [53, 54]. Importantly, learning a brain mapping to a closed set of acoustic features, such as spectrogram frequency features or spectrotemporal modulation features, can allow open-vocabulary speech decoding. While previous work has shown that acoustic information is a promising candidate for decoding [20, 21, 26, 27], it remains unclear to what extent acoustic properties of speech may be decoded directly from sensorimotor cortex using HD ECoG recordings and how intelligible resulting reconstructed speech would be.

A number of speech decoding models from brain activity have been proposed in the past. Recently, deep learning models have become increasingly more popular in the BCI field due to their potential for learning complex relationships between sources and targets of decoding [18, 20, 24–26, 33, 55]. However, there is a plethora of deep learning architectures [56] and parameter choices. As a result, many studies tend to borrow model architectures and parameters from other fields and apply them with minimal changes to neural data for speech decoding. Thus far, no comprehensive study on optimization of deep learning models for speech reconstruction has been performed. Moreover, there is a lack of consensus regarding choices of brain and audio speech features that are used in such models. Current studies use different methods for ECoG data preprocessing, including denoising reference schemes and frequency component ranges, as well as different types of audio speech features. Finally, evaluation of speech decoding models is also approached differently, and consensus regarding which metrics may be more informative is currently lacking. Overall, it remains unclear whether such a neuroengineering approach for speech reconstruction from the motor cortex could provide the basis for a real-world autonomous BCI application for decoding intended speech, and what the next practical steps for it would be.

In the present work, we aimed to decode acoustic properties of speech from sensorimotor cortex by building and evaluating deep learning models for the reconstruction of speech acoustics. For this, we collected HD ECoG data during a word production task, in which five human subjects spoke 12 unique words out loud, repeating each word ten times (figure 1). Having ruled out the possibility of acoustic contamination of neural signals (see section 5.6) [57], we performed a dedicated machine learning procedure to optimize parameters of speech reconstruction models and examined metrics for evaluation of the reconstruction results (figure 2). First, we report a significant improvement



of the reconstruction accuracy due to model optimization and identify parameters associated with best results. Second, we show that more complex decoding models lead to better speech reconstruction. Third, we found that individual metrics for evaluation of speech reconstruction reflect different aspects of model performance and that low- and high-level metrics can be dissociated. Fourth, decoding of individual words from speech reconstructed from sensorimotor brain activity achieved 92%–100% accuracy (chance level is 8%). Fifth, reconstructed speech exhibited high perceptual quality due to the model optimization procedure. Altogether, these results demonstrate the potential for reconstruction of identifiable, high-quality speech directly from sensorimotor brain activity and the importance of model optimization in achieving best perceptual results. These findings contribute to the state of the art in the field of speech decoding and reconstruction for BCI and have the potential to guide its further development.

2. Results

For speech reconstruction from sensorimotor brain activity, we optimized and evaluated three widely-used artificial neural network architectures: a multi-layered perceptron (MLP), a densenet convolutional neural network (DN) [58] and a sequence-to-sequence recurrent neural network (S2S [59], figure 2). First, we report results of the optimization procedures (how successful they were and what parameters of the reconstruction model were optimized). Next, we report low-level metrics for evaluation of the reconstruction accuracy (Pearson correlation, match in voice activity detection (VAD), short-term objective intelligibility (STOI)) and high-level metrics (identifiability of words and speakers with machine learning classifiers and speech perceptual intelligibility with human perceptual judgments). Finally, to determine which neural sources were most informative for achieving best results, we assessed the


contribution of individual intracranial electrodes to the reconstruction accuracy.

2.1. Model optimization

2.1.1. Optimization of the reconstruction loss in validation and test data

The three model architectures (MLP, DN and S2S) were optimized separately due to differences in their hyperparameters. Each model was optimized per subject using the Optuna (<https://optuna.org>) framework [60]. We used a form of Bayesian optimization that constructed a distribution of the reconstruction loss given various parameter choices by taking into account the history of parameter changes (see section 5 for details). Per HD ECoG subject, all task data was divided into train, validation and test subsets (figure 2). Per word, one randomly chosen repetition was added to the validation set (12 words in total) and another randomly chosen repetition was added to the test set (12 word in total, see section 5 for details). During training, the models minimized the reconstruction loss on the train set, measured as a mean squared error between target and reconstructed speech spectrograms. The reconstruction loss computed on the held-out validation set was used to optimize model parameters. After the optimization was complete, the reconstruction loss was computed on a separate test set to evaluate the results using leave-one-out cross-validation (LOO-CV), wherein, per each of 12 words in the test set, the model was retrained using optimal parameters on all remaining data.

Overall, we observed a considerable drop in the reconstruction loss on the validation set, which consisted of a single trial of every unique word (figure 3(a)). We also report substantial differences across subjects with respect to the minimal loss achieved (S1 and S5 reached much lower overall loss, especially in DN in S2S), and apparent interaction between subjects and models (S1 and S5 showed more distinct loss profiles across models compared to S2, S3



and S4, who showed similar profiles for MLP and S2S and whose minimal loss in DN and S2S was greater compared to S1 and S5).

Because the validation loss was used to guide model optimization (see section 5 for details), it was possible that optimal models became overfitted to the validation data and therefore would not generalize to a held-out test set. We tested this possibility with a Wilcoxon signed-rank test in all subjects using loss values in a test set obtained with non-optimized and optimized models (figure 3(b)). Non-optimized models used a parameter set that was chosen at random at the first step of the optimization process. The result showed a significant decrease of test loss in optimized models compared to non-optimized ones: $Z_{\text{all}} = 5.9$, $p = 2 \times 10^{-9}$, and separately per model: $Z_{\text{MLP}} = 3.15$, $p = 6 \times 10^{-4}$, $Z_{\text{DN}} = 5.01$, $p = 3 \times 10^{-7}$ and $Z_{\text{S2S}} = 1.93$, $p = .03$ based on non-parametric Wilcoxon tests. This means that the observed optimization effects generalized to unseen test data that had not been part of optimization.

2.1.2. Optimal parameters and parameter importance

Since optimization of speech reconstruction models overall led to a decrease of the reconstruction loss, we subsequently identified the optimal parameter configurations associated with best performing models. The optimization routine tuned three types of parameters: (1) input (brain data), (2) output (target mel-spectrogram), and (3) model-specific hyperparameters. All parameters and their ranges were selected based on previous literature and own experience. Parameters of the brain data (model inputs) included the type of reference for de-noising HD ECoG signals (common average reference [CAR] or bipolar), specific HFB frequency range of HD ECoG signals (60–300 or 70–170 Hz) and the temporal window of data used for modeling one spectrogram time point (160 or 360 ms). For reconstruction, we chose mel-spectrograms of speech as they represent sound acoustics on a spectrum inspired by human auditory perception [61]. Parameters of mel-spectrograms (model outputs) included the type




Figure 3. (a) Optimization history for the validation set plotted over the course of the optimization procedure, i.e. Optuna optimization ‘trials’ (see section 5 for details). (b) Optimization effects in the test set. Reconstruction loss computed on the test set is shown for non-optimized (a parameter set chosen at random at the start of the optimization procedure, i.e. first Optuna ‘trial’) and optimized models. (c) Optimal parameters per model architecture (MLP, DN, S2S) and subject. Different colors represent different types of parameters: input (purple), output (red), general model hyperparameters, such as dropout and learning rate (gray) and model-specific parameters: MLP (blue), DN (orange) and S2S (green). These plots represent values of each parameter in the optimal model. For example, for S1 MLP best parameters were: CAR, 60–300 Hz, 160 ms window—input; 40 features of clean audio at 75 Hz—output; dropout of $2e-4$, learning rate of $1e-4$, four layers of 256 nodes each. (d) Parameter importances per model. Importance scores were normalized across parameters and are shown per model-subject pair.

of audio preprocessing: microphone audio as it was acquired (raw) or de-noised microphone audio that filtered out all acoustics that were not produced by the subject speaking (clean, see section 5 for details), and details of spectrogram computation: the number of mel-frequency bins (40 or 80) and the sampling rate (15, 25 or 75 Hz, see section 5 for details). Finally, model-specific hyperparameters included common for all models (learning rate and dropout) and architecture-specific parameters (number of nodes and number of layers in MLP; bottleneck, growth and reduce factor, number of blocks in DN, number of encoding and decoding layers and bidirectionality in S2S).

To explore optimal parameters, first, we visualized optimal parameter configurations for each model-subject pair and each type of parameter (figure 3(c)). Inspection of these plots allowed us to spot consistencies across models and subjects, such as dominance of CAR, 60–300 Hz of high gamma range and 360 ms temporal window among the optimal configurations of the input; preference for 40 mel features in clean audio at lower sampling rates (15 or 25 Hz) in the optimal configurations of the output. Across all model architectures there appeared to be a preference for a larger number of layers and nodes. In the case of S2S, optimal models in all subjects included bidirectionality.




Figure 4. (a) Examples of original and reconstructed spectrograms (test data, DN model, S5). (b)–(d) Low-level evaluation metrics computed on test data and shown for non-optimized and optimized versions of each model: Pearson correlation (b), voice activity detection (VAD) match (c), short-term objective intelligibility (STOI) (d). Boxplots outline 25th and 75th percentile of results across test trials of all subjects. Red line shows medians. Semi-transparent smaller markers show results per individual test trial of each subject. Opaque bigger markers show averages across test trials per subject. Lower (based on data shifts) and upper (based on target audio waveforms) performance bounds are shown in gray and pink, respectively. Results for the three model architectures are shown in different colors (MLP in blue, DN in orange and S2S in green). Color saturation denotes the use of optimization: results for non-optimized models are shown as non-saturated and results of optimized models are saturated. Significant differences in medians are marked: $p < .05$ (*), $p < .01$ (**) and $p < .001$ (***).

To examine the effects that different parameters had on the loss values, we used parameter importance estimates provided within Optuna. Parameter importance was estimated based on a non-linear regression that used parameter values and their interactions as predictors of the reconstructions loss values. Larger importance was associated with larger prediction strength of specific parameter combinations. Given the large number of parameters and their interactions, here, we only interpreted parameters with the largest importance levels. Overall, across subjects, we observed that the reconstruction loss was most affected by output parameters, such as type of audio and sampling rate (figure 3(d)). For MLP, the number of mel features and number of layers were the next most important parameters. For DN, the size of the input temporal window, dropout and learning rate were the next most important parameters. For S2S, the number of encoding layers, learning rate and dropout were the next most important parameters.

These results indicate that optimal configurations of parameters lead to lower reconstruction losses. Specifics of audio processing may be of particular importance, while the exact values of other

parameters may be less relevant given their lower importance. Optimal parameters were often consistent across models and subjects. Yet, it also appeared that some parameters may be subject- and model-specific, highlighting the need for subject- and model-specific optimization for achieving best reconstruction results.

2.2. Evaluation with low-level measures

Having reached optimal model performance in terms of the reconstruction loss, next, we sought to evaluate the reconstruction quality. After a brief visual inspection of the reconstructed and target spectrograms (figure 4(a)), we turned to objective metrics of reconstruction quality typically used in the literature [19–21, 26]. These are measures based on low-level acoustic features that are easy to calculate and interpret. Three metrics were considered: (1) Pearson correlation between reconstructed and target speech fragments, (2) match in VAD between reconstructed and target fragments, and (3) STOI measure [62]. VAD match calculates the correspondence in detected voice activity between reconstructed and target waveforms. STOI, originally developed for comparison of clean and degraded signals, is an estimate of

speech intelligibility based on the spectrotemporal feature correspondence between the audio fragments. It is often used to assess intelligibility of speech reconstructed from brain activity [19–21, 63]. In addition, we provided lower (based on random shifts) and upper (based on target audio fragments) bounds of model performance. We also focused on two comparisons of metric values: (1) between optimized and non-optimized models (within-model), and (2) across the three optimized models (MLP, DN and S2S).

Median Pearson correlations across subjects, calculated between reconstructed and target spectrograms, were significant against chance for all models: $p = 3 \times 10^{-3}$, $p = 8 \times 10^{-4}$, $p = 0$ for non-optimized MLP, DN and S2S, respectively, and $p = 1 \times 10^{-3}$, $p = 0$, $p = 0$ for optimized MLP, DN and S2S, respectively, based on a non-parametric permutation test. We found the largest optimization effects compared to VAD match and STOI (figure 4). The difference between optimized and non-optimized reconstruction models in terms of Pearson correlation scores was significant for MLP and DN only: $Z_{\text{MLP}} = 2.67$, $p = 4 \times 10^{-3}$, $Z_{\text{DN}} = 5.85$, $p = 2 \times 10^{-9}$ and $Z_{\text{S2S}} = 1.66$, $p = .05$ based on non-parametric Wilcoxon tests with Bonferroni correction. This was somewhat expected since Pearson correlations are related to the reconstruction loss (measured as the mean squared error between predictions and targets). In addition, Pearson correlations in DN were significantly greater compared to MLP and S2S: $H = 14.26$, $p = 8 \times 10^{-4}$, based on a non-parametric Kruskal–Wallis test and post-hoc Dunn test on individual comparisons: $p_{\text{DN-MLP}} = 1 \times 10^{-4}$ and $p_{\text{DN-S2S}} = .04$, corrected for pairwise multiple comparisons.

VAD match and STOI were calculated on synthesized waveforms. This allowed us to provide ceiling estimates for each subject-model pair by computing the metric value on original audio fragments (used as input to the reconstruction models) and audio fragments resynthesized from spectrograms to waveform using the external speech synthesis model (Parallel WaveGAN vocoder). Median VAD match values were significant against chance for all models except for non-optimized MLP: $p = .06$, $p = 0$, $p = 0$ for non-optimized MLP, DN and S2S, respectively, and $p = .03$, $p = 0$, $p = 0$ for optimized MLP, DN and S2S, respectively, based on a non-parametric permutation test. Median STOI values were only marginally significant against chance: $p = .02$, $p = .03$, $p = .02$ for non-optimized MLP, DN and S2S, respectively, and $p = 5 \times 10^{-3}$, $p = 2 \times 10^{-3}$, $p = .11$ for optimized MLP, DN and S2S, respectively, based on a non-parametric permutation test. Only DN showed an effect of model optimization for STOI: $Z_{\text{DN}} = 3.27$, $p = 5 \times 10^{-4}$ based on non-parametric Wilcoxon tests with Bonferroni correction. There was no significant effect of model architecture.

Similar to the results from the section on model optimization, we also observed considerable variance in metric scores across five subjects with S1, S2 and S5 on average obtaining greater Pearson correlation and STOI scores compared to S3 and S4.

It is important to note that these metrics rely on low-level acoustic features and are convenient to consider when evaluating audio reconstruction quality (especially given the low-level spectrogram feature loss used in training). In the light of the ultimate goal of using these reconstruction models in clinical BCI applications, it does not suffice to limit evaluation to low-level quality metrics. Therefore, next we assess reconstruction of high-level characteristics of speech.

2.3. Evaluation with high-level measures

We analyzed high-level properties of the reconstructions, such as identifiability of individual word and speakers, as well as intelligibility and perceptual quality of resulting speech. First, we assessed word and speaker identifiability with objective metrics, such as machine learning classifiers. Then, in two separate speech perception experiments, we collected subjective human judgments of word intelligibility and speaker recognition. Finally, in a third behavioral experiment, we collected and analyzed perceptual quality judgments of the reconstructions.

2.3.1. Machine learning classifiers

We trained a linear logistic regression classifier to decode word identity (out of 12 individual words). For audio reconstructions, a linear logistic regression classifier was trained on target audio spectrograms per subject and tested on corresponding reconstructions from the test set per subject and model (including optimized and non-optimized versions of each model). Word classification accuracy was highly significant (reaching accuracy of 92%–100%, chance level of 8%) across all subjects and models. There was no significant difference in accuracy between optimized and non-optimized models, different model architectures or subjects (figure 5(a)). In addition, we examined classifiers' confidence about the predicted word classes expressed in probability distributions over all candidate classes per trial. The median confidence was 1 suggesting that word classification from speech reconstruction was robust (figure 5(b)). To test this further, we performed a noise study, in which we repeatedly added white noise (from standard normal distribution) to the input training and test data of the classifier, retrained it and recalculated its accuracy on the test data. We observed that word identification accuracy did not change as a function of amount of noise, resulting in robust and stable classification accuracy, even when considerable amounts of noise were added (figure 5(c)).

For reference, we retrained each logistic regression on brain input rather than audio spectrograms




Figure 5. (a) Word identifiability as assessed with word classification using audio reconstructions from optimized and non-optimized versions of the three models (MLP, DN, S2S). The three model architectures are shown in different colors (MLP in blue, DN in orange and S2S in green). For reference, word classification using raw brain input (to be passed as input to the speech reconstruction model) is also shown. Color saturation denotes the type of input data: brain input (least saturated), speech reconstructions obtained with non-optimized models (medium saturated) and speech reconstructions obtained with optimized models (most saturated). In the case of brain inputs, brain data with optimal input parameters per model (sampling rate, reference and HFB range) was used. (b) Probability matrices for predicted and target test words for a classifier trained on audio spectrograms and tested on reconstructions. Top panel shows probability matrices averaged over subjects. Optimized DN model was used to obtain the reconstructions. The lower panels represent matrices of best (S1) and worst performing (S3) subjects with respect to the word classification accuracy. For reference, next to the matrices computed using audio reconstructions with optimized models, matrices computed using raw brain input are shown. (c) Results of the noise study: performance of the classifiers trained and tested on noisy inputs. The amount of noise was gradually increased. Individual markers show performance per subject. Results from all three models are aggregated together per bar. Individual subject markers reflect performance of individual models (MLP in blue, DN in orange and S2S in green). For reference, results using raw brain input are shown in pale purple. Results using optimized reconstructions with optimized models are shown in dark purple. (d) Speaker identifiability as assessed with speaker classification using classifiers trained on audio spectrograms and tested on reconstructions. Results for optimized and non-optimized versions of the three models (MLP, DN, S2S) are shown. A confusion matrix for the best performing classifier (optimized DN) is shown on the right.

and tested it on the same test set of 12 words as above. Interestingly, we observed larger variance in word classification accuracy across subjects compared to the use of audio spectrogram reconstructions (figure 5(a)), consistent with inter-subject variability patterns from low-level evaluation metrics (lower performance of S3 and S4). Median classifier confidence using brain input was .3 and the accuracy decayed rapidly as more noise was added to the brain input. It is important to note, however, that using raw brain input in word classification with logistic regressions is at a disadvantage compared to using speech reconstructions as it does not include non-linear transformations of brain signals that optimized reconstruction models provide.

Next, we trained logistic regression classifiers to decode speaker identity. As with word classification, the classifiers were trained on target spectrograms and tested on reconstructions for optimized and non-optimized versions of each model. We observed the highest accuracy for optimized DN model, reaching 73% ($p_{DN} = 0$ based on random label permutations,

chance level is 20%, figure 5(d)). There was a trend for optimized models to perform better (in DN and MLP, but not S2S). DN and S2S models achieved higher accuracy compared to MLP. The latter provided insignificant decoding accuracy: $p_{MLP} = .45$, $p_{MLP} = .2$ for non-optimized and optimized models, respectively. Due to the difference in the number of channels in input brain data, no classifiers on brain input were trained for speaker decoding.

2.3.2. Human perceptual judgments

With follow-up speech perception experiments, we sought to test whether word and speaker recognition not only were successful with machine learning classifiers trained on target spectrograms, but were also possible with human perceptual judgments.

We conducted three experiments: a word recognition experiment (I), a speaker recognition experiment (II) and an audio comparison experiment (III, see section 5 for details). In experiments I and II, 30 and 29 healthy Dutch-native volunteers respectively, were presented with an audio reconstruction




Figure 6. Human perceptual judgments obtained with three behavioral experiments. (a) Word intelligibility (experiment I). (b) Speaker recognition (experiment II). (c) Perceptual quality of optimized and non-optimized model reconstructions (experiment III). Perceptual judgments per each experiment were plotted against low-level evaluation metrics (scatter plots to the right of each boxplot panel). Boxplots outline 25th and 75th percentile of behavioral results. Red line shows medians. Semi-transparent smaller markers show results per behavioral participant on data of each HD ECoG subject (averaged over 12 test word trials). Opaque bigger markers show averages across behavioral participants per HD ECoG subject. Results for the three model architectures are shown in different colors (MLP in blue, DN in orange and S2S in green). Lower (based on randomly assigned perceptual judgments) and upper (based on target audio waveforms) performance bounds are shown in gray and pink, respectively. Experiment III did not use target audio waveforms and therefore does not have the upper performance bound. Significant differences in medians are marked: $p < .05$ (*), $p < .01$ (**) and $p < .001$ (***). Boxplots in experiments I and II show accuracy of word and speaker recognition based on binary comparisons. Boxplots in experiment III show whether reconstructions from optimized or non-optimized models were judged to be better (more intelligible, less noisy). An arbitrary value (perceptual quality weight) of 1 was assigned to each judgment if the audio produced by the optimized model had been judged as better and a value of -1 otherwise (see section 5 for more details). In order to make scatter plots of perceptual judgments against low-level metrics, perceptual judgments were recalculated per test word trial (averaged over behavioral participants). Therefore, each data point of a scatter plot represents data per single word from the test set of each HD ECoG subject – 60 data points in total (12 word trials $\times$ 5 subjects).

and two options, of which they were asked to choose one that suited the reconstruction best. In the word intelligibility experiment, the two options were two written words (for example, ‘janneke’ and ‘grootmoeder’), one being the correct word, and the other chosen at random from remaining 11 words. In the speaker recognition experiment, two options were target audio fragments of two speakers pronouncing the same word as in the reconstruction (for example, ‘janneke’ said by S1 and S4), one being the correct speaker, and the other chosen at random from remaining four speakers. Both experiments included catch trials that used target audio fragments instead of reconstructions. This was done to obtain ceiling estimates of word intelligibility and speaker recognition.

In the case of word intelligibility, participants on average were able to identify words significantly above chance across subjects and models (except for some models of S3 and all models of S4): median accuracy values reached 58%, 58%, 67%

for MLP, DN and S2S models, respectively, all significant against chance ($p = 4 \times 10^{-4}$, $p = 1 \times 10^{-3}$ and $p = 0$, respectively, based on a non-parametric permutation test). Best scores were obtained for DN and S2S models: $H = 7.67$, $p = .02$, based on a non-parametric Kruskal–Wallis test and post-hoc Dunn test corrected for pairwise multiple comparisons: $p_{S2S-MLP} = 7 \times 10^{-3}$ and $p_{S2S-DN} = .06$. (figure 6(a)). Moreover, there was a positive relationship between low-level Pearson correlation and high-level word intelligibility metric for DN and S2S models.

In the speaker recognition task, participants showed even better overall performance and a large gap between behavioral scores and permutation-based lower bound: median accuracy values reached 75% for all models, which was significant against chance (all p -values are equal to zero based on a non-parametric permutation test). There was no significant difference in scores across models and the only model with a positive relationship between Pearson

correlation and speaker recognition judgments was MLP (figure 6(b)).

These results differ from the performance of machine learning classifiers for both word intelligibility and speaker recognition. In the case of word intelligibility, machine learning classifiers reached an accuracy of 92%–100% (chance level 8%) in identifying individual words, whereas behavioral judgments indicated that perceptual quality of reconstructed words was considerably lower and provided an average accuracy of 58%–66% (chance level 50%). In both cases, the MLP model was outperformed by DN and S2S. For speaker recognition, machine learning classifiers showed best results for DN and S2S as well: 73% and 57% respectively (chance level 20%), with behavioral judgments achieving 75% (chance level 50%) for all models. Due to the differences in the computation of accuracy between machine learning classifiers and perceptual judgments, direct comparison between the two results cannot be performed. Moreover, behavioral experiments were performed on synthesized speech waveforms whereas machine learning classifiers used normalized spectrogram data. In order to synthesize waveforms the spectrograms were re-scaled and centered with mean and standard deviation of the target audio and passed through the external vocoder. This was necessary for correct processing by the vocoder and is a standard step for processing new data by a pretrained model [64].

In another behavioral experiment (experiment III), we sought to assess model optimization effects on perceptual scores. Similar to the previous experiments, 29 participants were presented with two options: non-optimized and optimized reconstruction audio fragments, otherwise matching in word, speaker and model (for example, optimized and non-optimized ‘janneke’ reconstruction using the MLP model on data of S1). No catch trials using target audio were used. We found that, on average, participants judged reconstructions from optimized models to be perceptually better across all models. All median values across subjects were positive: .17, .25 and .17 for MLP, DN and S2S models, respectively, all significant against chance ($p = 6 \times 10^{-4}$, $p = 0$ and $p = 8 \times 10^{-4}$, respectively, based on a non-parametric permutation test). Largest perceptual gains were obtained for optimized DN and MLP models: $H = 18.74$, $p = 8 \times 10^{-5}$, based on a non-parametric Kruskal–Wallis test and post-hoc Dunn test corrected for pairwise multiple comparisons: $p_{\text{DN-S2S}} = 3 \times 10^{-5}$ and $p_{\text{MLP-S2S}} = 3 \times 10^{-3}$ (figure 6(c)). This result confirmed not only that optimization led to a lower reconstruction loss and more accurately preserved low-level features (measured with Pearson correlation), but also that optimized models were associated with better perceptual quality of the reconstructions.

Altogether, our behavioral results further supported the previous conclusions made based on the loss analysis, low-level and objective high-level evaluation metrics, and confirmed overall high-quality of the optimized speech reconstructions.


2.4. Contribution of individual electrodes

Finally, given overall high-quality reconstruction accuracy, as assessed with multiple low-level and high-level metrics, we sought to investigate the relationship between model performance and input brain features. Understanding interactions between electrode location on the sensorimotor cortex and its influence on speech reconstruction quality is vital for future applications of this work in the BCI field. To assess this, we performed an input perturbation analysis to assess individual contribution of each HD ECoG electrode to the reconstruction accuracy. For this, during testing (after models had been normally trained) per electrode we perturbed brain data that was input to the reconstruction models (optimized MLP, DN or S2S) by replacing that electrode’s time course with zero values (done after centering the input data) and recording the resulting loss value. Thus, we were able to associate perturbation of data per electrode with the amount of change in the loss value compared to using all input data.

Such input perturbation analysis revealed that most informative electrodes (those associated with largest increases in the loss value, and consequently, worse reconstruction accuracy) were grouped in small clusters along the ventral and dorsal premotor and motor areas, rather than being evenly distributed across the HD ECoG grids (figure 7). We also observed a degree of variability across patients and, to a smaller extent, across reconstruction models. This, together with previously mentioned differences across subjects may indicate that different cortical locations can be more or less informative for speech decoding across subjects, and that an individual subject approach is beneficial for achieving high accuracy of results.

3. Discussion

In this study, we performed a systematic optimization and evaluation of models for speech reconstruction directly from sensorimotor brain activity recorded with HD ECoG grids. We found that end-to-end deep learning approaches, currently dominating the field, overall benefited from model optimization, and that the choice of output parameters of these models (target spectrograms) had the largest effect on the reconstruction quality. Next, we showed that optimized models led to fairly good reconstruction quality as assessed with both low-level evaluation metrics and high-level speech measures, such as word and speaker



identifiability and perceptual quality. Word recognition in reconstructed audio was markedly more accurate, stable across subjects and robust compared to raw brain input. Behavioral experiments showed certain individual word intelligibility and optimization gains in perceptual judgments. Finally, we quantified the relationship between the reconstruction accuracy and location of HD ECoG electrodes, revealing that the largest contribution to model performance was made by small clusters of electrodes throughout the ventral and dorsal premotor and motor cortices. These results have the potential to further advance the state-of-the-art in speech decoding and reconstruction for subsequent use in BCIs for communication in individuals with severe motor impairments.

3.1. Optimized complex model architectures lead to best reconstruction of spoken speech from brain signals

For all models, optimization led to a decrease in the reconstruction error. Our results indicate that input parameters, or options for preprocessing brain

data, and model-specific hyperparameters may affect model performance compared to the choices of output audio parameters. Specifically, the quality of microphone recordings and audio denoising procedures seem to be key for achieving best performance. This appears logical, but the clinical setting in which ECoG data are often acquired (especially for experiments with participants during awake surgery) can be a challenging environment for collecting high-quality data, and this needs to be taken into account when planning and conducting the research. Moreover, this result may contribute to the discussion about the type of information that drives neural responses. Background noise removed from ‘clean’ audio signals was part of the auditory input to the brain, yet in most subject-model pairs, including S1 who had limited auditory coverage, reconstruction was best if only produced speech, and not all perceived sound, was used. Based on these results, our recommendation for further BCI research on speech reconstruction in non-disabled participants is to ensure good-quality audio recordings, denoise microphone signals and use spectrogram features

over 40 frequency bins at the sampling rate of 15–25 Hz.

Next, we observed that computationally more complex models, such as a convolutional densenet (DN) model and a recurrent S2S model, outperformed a simpler MLP architecture on a number of performance metrics. These included Pearson correlation between reconstructed and target audio (DN), word and speaker classification accuracy (DN and S2S) and perceptual judgments of word identifiability (S2S). Moreover, MLP architectures with a larger number of layers (4 from the range of 1 to 5 layers) and nodes (208 and 256 from the range of 16 to 256 nodes) turned out to generate the best MLP results. These results suggest that high complexity of deep neural network models is required to achieve more accurate speech reconstructions. Our results also indicate larger optimization effects in DN and MLP compared to S2S (figures 3(b), 4(a) and 5(c)). And even though DN occasionally outperformed S2S, the improvements were marginal, which makes it difficult to conclude that convolutional architectures should be preferred over recurrent ones. It is important to note that computational complexity refers to the type of computed operations, and not the number of trainable parameters. Specifically, DN benefits from local convolutions and skip-connections, and S2S has state memory and an attention mechanism, whereas MLP only uses basic linear operations followed by a non-linear activation function.

Recent work in computational neuroscience suggests that the use of recurrent connections and attention modules can be beneficial in processing brain signals [65–67]. Recurrence can help capture local and long-range temporal dependencies in the data whereas attention helps to effectively process temporally organized inputs. It is important to note that the present results have been obtained with relatively small datasets, in which subjects pronounced individual words. It is possible that for speech data that exhibits longer temporal dependencies, such as speech phrases, sentences and narratives, attention-based encoder–decoder models will further outperform convolution-only architectures, such as DNs [26].

3.2. Individual evaluation metrics highlight different aspects of model performance

Evaluation of speech reconstructed from sensorimotor cortex was performed using various low- and high-level metrics. We observed that each metric assessed model performance differently. Pearson correlation metric showed the largest improvement of the reconstructed speech over the permutation-based baseline and was the only low-level metric to reflect

model optimization gains. In general, this result appears logical as there is a straightforward relationship between the reconstruction loss (mean squared error) and Pearson correlation of predictions and targets. In the future, more sophisticated loss functions may be explored especially for achieving higher perceptual quality of the reconstructions, which could be regarded as the ultimate test of performance in BCI applications for restoring communication. VAD match metric showed improvement compared to the baseline, yet no model optimization effects. It is possible that distinguishing between reconstructed speech and silence may have been a simple enough task that did not require optimization. Finally, due to the lack of significance and overall low values, STOI metric did not seem to be informative above and beyond what the other two measures showed. This could in part be due to the way it is calculated in that it requires audio fragments to be at least 300 ms in length. This limitation can be particularly inconvenient when reconstructing individual monosyllabic words. Altogether, low-level assessment of reconstruction results indicated that Pearson correlation remains the simplest and most straightforward low-level metric for evaluation of reconstruction accuracy in individual words, as long as the mean-squared error, otherwise known as the ‘pixel loss’, is used for model training. VAD and VAD match can remain useful in identifying speech fragments in the brain signal, especially if such identification is not trivial, such as during covert or attempted speech. STOI as an evaluation metric may potentially be phased out.

The present study also explored several high-level metrics. Machine learning classifiers demonstrated high levels of word and speaker identifiability in reconstructed speech. Perceptual judgments of speech intelligibility seemed to agree to some extent with Pearson correlation and machine learning classifiers. It was clear, however, that perceptual quality provided a unique measure that could not fully be captured in any other objective measure considered here. Therefore, we propose that future work relies less on commonly used metrics, such as Pearson correlation and STOI, and instead prioritizes human perceptual judgments of speech reconstruction quality.

3.3. Direct speech reconstruction from sensorimotor cortex as the basis for next-generation BCI applications

The results of this study indicate that in addition to fairly good audio reconstructions, the presented models can help achieve accurate word classification with a rather simple linear classifier. The reconstruction-based word decoding was highly

stable across subjects, resulted in decreased subject-specific variability (figure 5(a)), and was robust against noise (figures 5(b) and (c)).

It is not trivial to directly compare these results with other word decoding work from sensorimotor cortex in the BCI field. Previous work has shown significant word classification results during speech production by decoding from raw ECoG brain input (average accuracy of $85 \pm 13\%$ for 45 binary word classification tasks [2], average accuracy of 47.1% for decoding 50 words [33] and average accuracy of 54.2% for decoding 26 words [35]). This required a relatively large number of repetitions of each word (30 repetitions of 10 words [2], 34 repetitions of 50 words [33], hundreds of repetitions over days and task blocks of 26 words [35]). The present study demonstrates that high accuracy and robust decoding can be achieved on rather small datasets (10 repetitions of 12 words) if using speech reconstructions for classification, thereby highlighting the potential of this approach for further use in BCI.

Training word decoders on external speech data (for example, publicly available large audio corpora) and testing them on reconstructions from ECoG data from unseen subjects may be a potential next step for future application of this work in BCI. This may help leverage effects of transfer learning across various speakers and aid in extracting higher-level acoustic features for better generalization. Highly accurate ($>90\%$ accuracy) and reliable discrete decoding of several classes of words (compared to the state-of-the-art binary decoding [31]) will help BCI devices offer more degrees of freedom to their users and thereby provide a new state of the art in the field. Moreover, further increasing the number of decodable words to a set of 50 or 100 commonly used words can lead to the development of the next-generation BCI technology for communication [33].

The goal of the present study was to investigate the potential of speech reconstruction with HD ECoG over the sensorimotor cortex. As a result, this study provides the best strategies for speech reconstruction and decoding in terms of decoding models and its parameters, evaluation metrics and, to some extent, indications for electrode placement (see below). Importantly, the current work has been performed in non-disabled participants. Replication of direct speech reconstruction and decoding from the brain of BCI users, such as individuals with paralysis, is not trivial due to the lack of microphone audio data. We envision two possible solutions for this. First solution would involve using speech audio from external speakers as targets and training the BCI user to mimic the speech they heard. Second solution could be based on transfer models [68, 69] trained on speech from non-disabled participants to learn speaker-invariant

speech representations. In addition, external sound generation models [70, 71] can be employed to aid speech reconstruction.

3.4. Informative neural sources may be distributed across sensorimotor cortex

Apart from the neural targets and decoding model, the results of the present study also contribute to the discussion about the neural sources of decoding and the details of the acquisition technique. Specifically, the present study showed successful speech reconstruction and word identifiability using HD ECoG electrodes over the sensorimotor cortex—one of the most promising regions in BCI implants for communication. It is important to note that other regions, such as Heschl's gyrus and superior temporal cortex, were not the focus of this study, as their primary function is auditory processing of perceived speech. In the present work, we used a technique called input perturbation analysis to identify HD ECoG electrodes that were most informative for speech reconstruction and word decoding. Our results indicate that multiple regions throughout the motor, premotor and potentially inferior frontal gyrus contribute to accurate speech reconstruction. This is largely in line with previous work [4, 9, 10, 15]. Interestingly, subjects with best reconstruction results (S1, S2 and S5) benefited from contributions of electrodes located in the dorsal premotor and motor cortex (figure 7). The subject who did not have dorsal motor coverage (S4) showed contribution of electrodes in the inferior frontal gyrus, but their contribution scores were lowest compared to those of other subjects, who had electrodes located in dorsal premotor and motor cortex. These results may indicate the potential of dorsal premotor and motor regions in speech reconstruction and are in line with the previous reports on speech decoding from that region [25, 28, 33].

3.5. Inter-subject variability of results and concerns for BCI

We also observed that despite optimizing individual subject's datasets, there was a large variance in model performance across subjects. This inter-subject variability is a common outcome in speech decoding studies [12, 20, 23] and can stem from several factors. One key factor is the choice of the brain area that is covered with electrodes. Missing an important patch of cortex can dramatically reduce performance. Electrode placement, however, is influenced by clinical considerations such as orientation of the grids to accommodate the leads, anchor veins that make optimal positioning impossible, and lack of knowledge about which cortical patches are most informative. Other factors include patient alertness and motivation, medication and epileptic spikes that confound

the brain signals. In the present study, data from S4 was collected during awake surgery, where participant alertness may have been considerably reduced. Bigger datasets with larger numbers of participants would be beneficial for systematic investigation of these factors. Their optimization at data collection, processing and decoding stages may be key for further progress in the field.

Importantly, we also find that although low-level features and perceptual quality of reconstructions varied across participants, word identifiability in reconstructed speech remained accurate and robust across all subjects. These results highlight the possibility of using direct speech reconstruction as a tool for boosting word decoding from brain data.

3.6. Limitations

The present work has a number of limitations.

Speech reconstruction in the present study is performed on a relatively small dataset of twelve words repeated ten times per participant, and therefore our results (specifically, high decoding accuracy of 92%–100%) should be interpreted with some caution. Human HD ECoG recordings from temporarily implanted grids are quite costly and rare data to collect. It is rarely possible to obtain more data from word repetition tasks in one subject than what is presented here.

Given a wide range of options and a lack of consensus regarding the effects of parameter choices (input, output and model-specific parameters) in the literature, here, we limited ourselves to a few sets of parameters. Our choices were informed by existing literature and own experience, however, several improvements can be considered in the follow-up work. First, some recent work has demonstrated the potential of lower-frequency information and its coupling with HFB for speech processing and decoding [19, 28, 72]. Therefore, future work on speech reconstruction from sensorimotor cortex could explore the effects of adding lower-frequency ECoG components. Second, several of the input and output parameters in the present work were considered to be categorical variables during optimization (for example, specific frequency range of HFB and the temporal window of brain data used for modeling one spectrogram time point) with only a limited number of values used. This could be improved in the follow-up work in order to consider a more appropriate distribution over these parameter choices, include a larger range of options and achieve better optimization. Finally, some analyses that explored optimization effects were limited by high computational costs. Specifically, the comparison of optimized and

non-optimized models was only fully performed by comparing optimized models to a single random set of parameters (figure 3(b)). To compensate for that, for one subject, we performed a more thorough comparison using several random parametrizations that overall confirmed previously obtained results (figure S1).

Finally, the present study implements an input perturbation analysis in order to identify electrode contribution to reconstruction accuracy. Our implementation of this method has a number of limitations. Specifically, it does not take into account multivariate activity that is likely to be informative for reconstruction and instead simplifies the spatial relationship between individual electrodes. Similar to the points made above, electrode contribution and selection could potentially be implemented as part of the optimization procedure as well.

4. Conclusion

In this study, we performed dedicated optimization and evaluation of different deep learning models for speech reconstruction directly from intracranial sensorimotor neural activity during a speech production task. We showed that machine learning optimization of speech decoding pipelines and detailed evaluation of model performance metrics allow for achieving more accurate and interpretable reconstruction results, and improve our understanding of the brain signal and its relation to the speech features. Altogether, our results indicate that direct speech reconstruction from sensorimotor brain activity provides highly accurate and robust word decoding performance and overall intelligible speech, and therefore, has the potential to advance the state of the art in real-world applications of BCIs for communication in severely paralyzed individuals.

5. Methods

5.1. ECoG experiment

5.1.1. Participants

Four participants (S1, S2, S3, S5, age 36, 30, 51 and 24, respectively; three females) with medication-resistant epilepsy were admitted to the intensive epilepsy monitoring unit for diagnostic procedures after they underwent temporary implantation with subdural ECoG electrode grids to determine the source of seizures and test the possibility of surgical removal of the corresponding brain tissue. In addition to clinical procedures, participants gave written informed consent to participate in scientific research that accompanied ECoG recordings and could be conducted between clinical procedures. Participants also gave written informed consent for implantation of an

additional HD ECoG grid over their sensorimotor cortex for research purposes. One more participant (S4, age 21, male) underwent clinical awake surgery for brain tumor removal. During the surgery, a grid of HD ECoG electrodes was placed subdurally over the sensorimotor and inferior frontal cortices while the subject participated in functional mapping (up to 10 min) to guide tumor resection. The participant gave written informed consent to acquire and use their HD ECoG data from functional mapping tasks for research purposes. The study was approved by the Medical Ethical Committee of the University Medical Center Utrecht in accordance with the Declaration of Helsinki (2013).

5.1.2. Stimuli and task

Stimuli consisted of 12 unique Dutch words: 'waar', 'jip', 'ik', 'ga', 'boe', 'kom', 'hoed', 'spelen', 'plukken', 'allemaal', 'janneke', 'grootmoeder'. These words were selected from a Dutch children's book 'Jip and Janneke' to maximize the occurrence of specific phonemes (/k/, /p/, /a/ and /u/) used in a different task [15].

During the word production experiment, each participant was presented with a visual cue (target word on screen) and instructed to read aloud the word shown on screen. Each word was presented ten times. The order of words was randomized over participants, resulting in 120 word trials in total. Inter-trial interval duration was randomized and on average was equal to four seconds in S1, S2, S3 and S5: 4.08 ± 1.52 in S1, 3.93 ± 1.48 in S2, 4.11 ± 1.74 in S3, and 4.06 ± 1.34 in S5. Due to time constraints during awake surgery, inter-trial interval in S4 was reduced to 2.49 ± 0.36 . The four participants implanted with ECoG grids for epilepsy monitoring (S1, S2, S3, S5) were also presented with 'rest trials' ('-' on screen), during which they were instructed to be silent. The participant who underwent awake brain surgery (S4) was not presented with 'rest trials' due to time constraints. The experiment took 9.18 min in S1, 9.09 min in S2, 8.96 min in S3 and 9.51 min in S5. S4 completed the task in 6 min and 2 s.

5.1.3. Experimental procedures

All participants were implanted with HD ECoG grids over the sensorimotor cortex. The suspected pathological regions did not extend to the sensorimotor region covered by the HD grids. This was clinically confirmed after implantation. In all participants, except for S2, HD grids were implanted over the left hemisphere. The hemisphere for HD ECoG implantation was the same as for clinical ECoG grids, and that was determined based on the clinical need. HD grid configurations differed slightly across participants with S1 implanted with 128 contacts, 1.2 mm exposed

diameter, inter-electrode distance 4 mm, and S2, S3, S4 and S5 implanted with 128 contacts, 1 mm exposed diameter, inter-electrode distance 3 mm.

In the experiment with participants in the intensive epilepsy monitoring unit (S1, S2, S3 and S5), words were presented to the participants on a computer screen (21 in. in diagonal, at about 1 m distance) using presentation software (Neurobehavioral Systems). In the experiment with the participant in the operating room (S4), words were presented on a tablet using custom Python scripts. In both cases, HD ECoG data were acquired using the NeuroPort neural recording system at a sampling rate of 2000 Hz (Blackrock Microsystems). Synchronization of the task with neural recordings was achieved based on event codes sent from stimulus presentation software to the recording computer.

In addition to the neural data, audio recordings were acquired using an external microphone (Audio-Technica AT875R), connected directly to the neural recording system. This ensured synchronization of the microphone recordings with neural data acquisition. Microphone data were recorded at 30 000 Hz.

5.1.4. ECoG data processing

HD ECoG data were loaded in MATLAB using the NPMK toolbox. The data were preprocessed using custom software. First, all electrodes with noisy or flat signal (based on visual inspection) were excluded from further analyses (S1: none, S2: 5 electrodes, S3: 1 electrode, S4: 17 electrodes, S5: 1 electrode). Second, a notch filter was applied to the remaining electrodes to account for the effects of line noise (50 Hz and its harmonics). Third, common average (CAR) or bipolar referencing (using the next nearest channel) was applied. Fourth, data were transformed to the frequency domain using Gabor wavelet decomposition at either 60–300 Hz or 70–170 Hz in 1 Hz bins with decreasing window length (four wavelength full-width at half maximum). Fifth, HFB amplitude was obtained by averaging amplitudes for the entire range of extracted frequencies. Sixth, the resulting time series per electrode were log-transformed and downsampled to 100 Hz. In total, four preprocessed HD ECoG files were made per participant: bipolar in the range of 60–300 Hz, bipolar in the range of 70–170 Hz, CAR in the range of 60–300 Hz and CAR in the range of 70–170 Hz.

In order to provide statistics of baseline HD ECoG activity, during the task, per participant, we identified a continuous fragment of at least five seconds in length, during which the participant remained silent and at rest. Mean and standard deviation of baseline neural activity was calculated per electrode. These values were used for normalization of HD ECoG

task data during training of speech reconstruction models.

Electrode locations were coregistered to the anatomical MRI in native space using computer tomography scans [73, 74] and FreeSurfer [75].

5.1.5. Microphone data processing

Microphone data were loaded in MATLAB using the NPMK toolbox and resaved in the wav format at the original sampling rate of 30 000 Hz. Event codes from the stimulus presentation software provided information about onset of each trial and the target word per trial, but to extract precise timing of the pronunciation of each word, we used the PRAAT toolbox (<https://praat.org>) [76]. Using PRAAT we manually annotated each participant's microphone recording file (wav) with onsets and offsets of individual words.

Per participant, we also performed audio denoising procedures to obtain 'clean' versions of each original, or 'raw', microphone recording. Denoising was performed with Audacity. First, a high-pass filter was applied at 60 Hz. Then, a built-in noise removal tool was used first to estimate background noise from a fragment of three to five seconds in length, and then to subtract the estimated noise signal from the rest of the audio. This process was repeated four or five times to obtain cleaner background. Finally, additional manual corrections were applied to remove irregular noise, such as occasional talking of staff present in the room, sounds of medical equipment or other background noise. Both raw and clean audio files were used in speech reconstruction.

For training of speech reconstruction models, we extracted spectrogram features from raw and clean audio files of microphone recordings. We extracted log-mel spectrograms in the frequency range of 80 to 7600 Hz, as they best approximate human speech perception. For parameter optimization, we varied the following spectrogram extraction parameters: number of mel frequency bins (40 or 80), window length (1470, 882, 525) and hop size (1470, 882, 294). The choice of the latter two parameters determined the sampling rate of the resulting audio: 15, 25 or 75 Hz, respectively.

5.2. Speech reconstruction

5.2.1. Modeling approach

Brain data was used as input and speech spectrograms as output. Only data during speech pronunciation was used in model training and testing. Spectrograms of the microphone audio data, that the models are trained, tested and validated on, are referred to as target audio spectrograms. The speech reconstruction model was trained to solve a regression problem of mapping 2d brain inputs (electrodes by time) onto 1d audio outputs (frequency vector at one time point):

$$\mathbf{y} = f(\mathbf{X}) + \epsilon, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^m$ is a vector of spectrogram frequency values per one time point and m frequency bins, $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a matrix of associated neural recordings with n corresponding to the number of data points over time and p corresponding to the number of ECoG electrodes, $f(\cdot)$ is a non-linear transformation function and $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ is a noise term. The transformation function is learned by the reconstruction model based on the artificial neural network. The reconstruction model is trained end-to-end.

5.2.2. Reconstruction model architectures


Three commonly used artificial neural network architectures were employed as the reconstruction models: an MLP, a DN and an S2S recurrent neural model (figure 8). MLP is an artificial neural network that consists of layers of nodes that apply linear transformations to their inputs, followed by a non-linear activation function. In our configuration, each linear layer was followed by batch-normalization and leaky ReLU (negative slope = 0.25) activation function.

DN [58] is a convolutional neural network architecture with skip connections that aim to handle vanishing gradients. It consists of blocks of convolutional layers. In a block, the outputs of all preceding layers are directly connected as additional inputs to all subsequent layers via concatenation. To accommodate processing of growing inputs, each subsequent layer grows in number of nodes (channels) by a specified *growth factor*. To compress the model and downscale its complexity, transition blocks are inserted after each convolutional block. They reduce the amount of information currently processed by the network by a value of the *reduce* parameter. In addition, *bottleneck* layers can be used to reduce the number of network channels in a block (by using 1×1 convolutions).

S2S [59] is a recurrent neural network architecture with an encoder and a decoder components. The encoder processes temporal information in the input, its representation in the hidden state at the last time point is used as input to the decoder. Additionally, an attention mechanism is often employed for better performance. Here, we used a form of global attention [77] to weigh encoder outputs and produce a context vector that is used as input to the decoder. Both encoder and decoder components in our S2S are based on a gated recurrent units [78, 79].

5.2.3. Parameter optimization

Optimization was performed per model architecture and per subject using tree-structured Parzen estim-



ators (TPE) as implemented in the Optuna (<https://optuna.org>) package [60]. This enabled us to optimize choices of categorical variables, yet came with a drawback of having to use independent samples. The latter meant that hyperparameter values were selected independently without considering interactions between them. The parameter optimization procedure aims to identify a set of model parameters θ that minimize the objective function $f(\theta)$ —reconstruction loss in our case:

$$\theta^* = \arg \min_{\theta \in \Theta} f(\theta). \quad (2)$$

In order to identify optimal parameters, TPE estimate the posterior probability $p(y|\theta)$, where y is the validation reconstruction loss. This is done using the Bayes rule:

$$p(y|\theta) \propto p(\theta|y)p(y). \quad (3)$$

The probability $p(\theta|y)$ is estimated using a decision boundary based on a threshold y^* :

$$p(\theta|y) = \begin{cases} l(\theta) & \text{if } y < y^* \\ g(\theta) & \text{if } y \geq y^* \end{cases}. \quad (4)$$

Effectively, TPE estimates two probability distributions of model parameters: $l(\theta)$ associated with y values below the threshold and $g(\theta)$ for y values above the threshold. The expected improvement EI of the optimization algorithm is proportional to the ratio $l(\theta)$ to $g(\theta)$. EI is maximized by iteratively drawing samples from $l(\theta)$ and $g(\theta)$, calculating their ratio and

selecting a set of parameters θ that maximizes that ratio:

$$EI_{y^*} \propto \left(\gamma + \frac{g(\theta)}{l(\theta)}(1 - \gamma) \right)^{-1}. \quad (5)$$

Using the TPE framework as implemented in Optuna, we optimized several input and output parameters and model hyperparameters. Separate input files were used for referencing (bipolar or CAR) and band width (60–300 or 70–170 Hz). The window of input data for reconstructing one time point of the audio could also be adjusted (160 or 360 ms). Separate output files could also be used depending on the number of mel features (40 or 80) and audio type (clean or raw). Depending on the target audio sampling rate (15, 25 or 75), corresponding audio and neural files were used.

Model parameters included architecture-specific and general parameters. General parameters included learning rate of the optimizer and dropout ratio. Parameters specific to the model architecture included number of layers and nodes per layer in MLP, presence of bottleneck layer, growth factor, reduce factor and number of layers in DN and number of encoding layers, number of decoding layers and bidirectionality in S2S.

5.2.4. Training procedures

Data was divided into train, test and validation sets in a pseudorandom fashion. One trial per each word was selected for validation (12 words in total) and another trial for each word was selected for test (12 words

in total). Pseudorandomization ensured that no consecutive trials were used in validation or test. It also attempted to distribute validation and test trials such that they were well spread-out throughout the entire task.

Each model was trained to take in input brain data and reconstruct target audio spectrogram. To reconstruct one time point of the audio spectrogram, a temporal window of varying length (160 or 360 ms) around the target time point was selected in neural data as input to the model. Each model was trained for 500 epochs, after which the reconstruction loss was calculated on the validation set. The Adam optimizer [80] with fixed parameters except for the learning rate was used. All training and testing procedures were implemented in PyTorch (<https://pytorch.org>). The models were trained using a single graphics processing unit NVIDIA GeForce RTX 2080 Ti. The number of trainable parameters differed per neural network architecture and its specific configuration.

The described model training procedure was embedded in the model optimization study by Optuna. A separate study was made per model (MLP, DN, S2S) and subject (S1, S2, S3, S4, S5). Each Optuna study was set to run for 100 'trials' (not to confuse with a task trial, 'trials' in an Optuna study refer to a single model training procedure with a fixed set of parameters). In each Optuna trial, a set of model parameters was selected, the model was trained for 500 epochs and the reconstruction loss on the validation set was calculated. The latter was used to guide parameter selection in the subsequent Optuna trials. In the end, the best Optuna trial with the corresponding parameter set was identified as the Optuna trial with the lowest reconstruction loss on the validation set.

Given the limited amount of data in this study, several techniques were employed to reduce overfitting. First, the standard dropout [81] technique was used. Second, a separate validation set was used to guide model and parameter choice. No data augmentation was applied.

5.2.5. Parameter importance

After model optimization, functional ANOVA (fANOVA) [82] as implemented in Optuna, was used to fit a random forest regression model and predict the reconstruction loss in the validation set given a specific parameter set. FANOVA estimated the effects of main variables (parameters) and interactions between them on the variance of the dependent variable (reconstruction loss).

5.2.6. Testing procedures

To obtain reconstruction results on the test dataset, LOO-CV was employed. LOO-CV refers to leaving one test word (test set consisted of 12 unique words) out of training each time. A separate model

was trained for 500 epochs per each word in the test set using the set of optimal parameters identified during model optimization. Data from the validation set was added to the training set to increase the amount of available data.

The model at the 500th epoch was used to obtain reconstructed spectrograms of the target test word. For all evaluation procedures, reconstructed and target spectrograms were computed per test word trial beginning with word onset and ending with word onset plus one second, which matched the duration of the longest trial across all subjects. An external vocoder, Parallel WaveGAN [64] (as implemented here: <https://github.com/kan-bayashi/ParallelWaveGAN>), was used to synthesise speech waveforms from the reconstructed spectrograms. Parallel WaveGAN is a generative deep neural network trained to synthesise speech waveforms given a mel-spectrogram. We used publicly available weights of the vocoder model (ljspeech_parallel_wavagan.v1) pre-trained on the LJ Speech dataset of spoken English (LibriVox project). Our experiments showed that despite being trained on English speech the model synthesized high-quality intelligible Dutch speech as well. In order to synthesise waveforms both target and reconstruction spectrograms were upsampled to match input dimensions the vocoder model was trained on: 80 mel-frequency bins at a sampling rate of 86.13 Hz. Each evaluation metric (except for speaker classification and human perceptual judgments) was then computed per LOO-CV fold (per one of the 12 word trials in the test set).

5.3. Low-level evaluation metrics

Pearson correlation, as implemented in the Scipy (<https://scipy.org>) library for Python, was computed between reconstructed and target spectrograms by vectorizing the spectrograms per word and correlating resulting vectors. This was done to account for the relative differences in spectrogram energy over frequency bins and time points.

Match in VAD and STOI [62] were computed on waveforms synthesized from reconstructed and target spectrograms using the WaveGAN vocoder. VAD was applied separately to each audio file (reconstructed and target speech). It was computed per each 30 millisecond window of audio and resulted in a string of 'one' (voice activity detected) and zero (no voice activity detected) values per file. We used python-based interface (<https://github.com/wiseman/py-webrtcvad>) for online free VAD service <https://webrtc.org>. VAD match was then calculated as the fraction of VAD output matches between reconstructed and target speech fragments:

$$match = \frac{\sum_{i=1}^n f^{(i)}}{n}, \text{ where } f^{(i)} = \begin{cases} 1 & \text{if } VAD_y^{(i)} = VAD_{\hat{y}}^{(i)} \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where y is the target waveform, $\hat{y}$ is the reconstructed waveform of equal length and n is the number of 30 millisecond windows in each file, per which one VAD value was computed.

STOI was computed on synthesized waveforms using a STOI package for Python (<https://github.com/mpariente/pystoi>). For those waveforms with a duration that was shorter than the minimum required for STOI computation (30 frames, or 300 ms, recommended for intermediate intelligibility), the algorithm set the output value to 10^{-5} . All such trials were excluded from subsequent analyses.

For each low-level evaluation metric, surrogate baseline distributions were computed for calculation of statistical p -values. This was done by shifting the neural data to time points corresponding to non-speech fragments 1000 times prior to obtaining test reconstructions and recalculating each low-level metric for each shift. This was done because one of the metrics, namely VAD match, only distinguished between speech and non-speech fragments and did not discriminate between individual words. Additionally, to provide another baseline for Pearson correlation and STOI metrics that focused on distinctions between individual words, we computed statistical p -values based on word permutations. For this, we used 1000 permutations of word labels when selecting data from the target audio files, thereby disrupting the correspondence between reconstructed and target speech fragments. The statistical p -values were comparable to those obtained with the shift-based baseline.

5.4. High-level evaluation metrics

5.4.1. Machine learning classifiers

Word classification was performed using logistic regression. Per each test word (out of 12, in the LOO-CV), the classifier was trained either on raw brain input (same data as input to the speech reconstruction models), or target audio spectrograms from the train set data. The classifiers were then tested on the test data: raw brain data or reconstructed spectrograms (obtained by passing raw brain input through the speech reconstruction models). To investigate word length effects, we additionally retrained the classifiers using audio fragments padded with Gaussian noise. This was done for all shorter words following guidelines from previous work [23]. The performance difference between original and noise-padded classifiers was not significant in any of the subjects: $Z_{\text{orig-noise}} = -0.03, p = 0.52$ for non-optimized models and $Z_{\text{orig-noise}} = 1.27, p = 0.1$ for optimized models based on non-parametric Wilcoxon tests.

In addition to this, a separate noise study was performed to estimate the performance robustness of word classifiers. We iteratively added incremental

amounts of white noise to the classifier input data during training and recomputed their performance on test data. In the case of the reconstructed data, the noise was applied after the reconstructions were obtained with the deep learning models.

Next, a speaker classifier was trained on target speech spectrograms of all participants. Per speech reconstruction model, it was tested on the reconstructed spectrograms. Python library Scikit-learn (<https://scikit-learn.org>) was used for training and testing all classifiers.

For both metrics, permutation tests (1000 shuffles of classification labels) were conducted to provide a surrogate baseline distribution for calculation of statistical p -values for each observed metric.

5.4.2. Human perceptual judgments

Human perceptual judgments of reconstructed spectrogram were collected from healthy volunteers in a series of online behavioral experiments. We conducted three experiments: a word recognition experiment (I), a speaker recognition experiment (II) and an audio comparison experiment (III). Participants (native Dutch with normal hearing; I: 30 participants, age 27 ± 8 , 12 females; II: 29 participants, age 27 ± 8 , 7 females; III: 29 participants, age 27 ± 8 , 8 females) were recruited online via Prolific (<https://prolific.co>). Each experiment was listed separately, and Prolific users were free to choose to participate in any number of experiments in any order. All participants gave their written informed consent and were reimbursed for their participation. The study was approved by the ethical committee of the Faculty of Social Sciences at Radboud University.

Experiments were implemented using Gorilla Experiment Builder (<https://gorilla.sc>). Because target and reconstructed microphone recordings were used in the experiments, to make voices of individual participants (S1, S2, S3, S4 and S5) unrecognizable, pitch of all audio files was shifted by 1 or 2 tones up or down, depending on the ECoG participant. In each experiment, behavioral study participants were told that the study investigated the effect of audio degradation on speech intelligibility. In experiment I, participants were presented with an audio fragment and two written words, and were instructed to select the word that corresponded to the audio they heard. In experiment II, participants were presented with three audio fragments: the 'target' audio at the top of the screen, and two audio fragments at the bottom, labeled 'speaker X' and 'speaker Y'. Participants were instructed to select the speaker label that corresponded to the speaker of the 'target' audio. In experiment III, participants were presented with two audio fragments and were instructed to select the fragment

with an overall better audio quality (more intelligible, less noisy). Data from each experiment, as well as screenshots of each experiment's layout have been made publicly available (see section 5.7). In experiments I and II, synthesized audio waveforms from both reconstructed (MLP, DN and S2S) and target spectrograms were used. Only optimized models were used for producing reconstructions, resulting in 360 trials: data from 5 ECoG subjects $\times$ 12 words from the test set $\times$ 3 speech reconstruction models $\times$ 2 types of audio files (reconstructions and targets). In experiment III only reconstructed audio were used for direct comparison of perceptual quality of MLP, DN and S2S outputs. In addition, both optimized and non-optimized (that used a randomly chosen parameter set from the first Optuna 'trial') models were used for producing and comparing the reconstructions, resulting in 360 trials: data from 5 ECoG subjects $\times$ 12 words from the test set $\times$ 3 speech reconstruction models $\times$ 2 types of optimization (optimized and non-optimized).

Behavioral judgments from experiments I and II were loaded into Python, response accuracy was averaged over test words, and the results were plotted per behavioral participant, HD ECoG subject and model (figures 6(a) and (b)). Responses from experiment III were assigned an arbitrary weight of 1 if reconstruction with the optimized model was rated as having higher quality, and a weight of -1 otherwise. The weights were then averaged over test word trials and plotted per behavioral participant, HD ECoG subject and model. Non-parametric tests were used for calculation of significance. For this, on every trial, instead of taking the participant's response, the choice was made at random. This procedure was repeated 1000 times. The surrogate distributions were used in calculation of statistical p -values for the observed human judgments.

Finally, we made scatter plots and fitted a regression line between results obtained with low-level evaluation metrics (Pearson correlation, VAD match and STOI) and human perceptual judgments

5.5. Input perturbation analysis

To estimate individual electrode contributions to the performance of the reconstruction model, an input perturbation analysis was performed. For this, per subject and model, we used the reconstruction loss of the optimized trained model on test data as our reconstruction loss baseline. Then, we iteratively recomputed the reconstruction loss on test data the number of times equal to the number of electrodes. Each time, prior to computing the loss but after input normalization, we zeroed out input of one of the electrodes. Then, we calculated the difference between the resulting reconstruction loss compared to the

baseline loss value (when input data from all electrodes was used). This way, if an electrode contributed to the reconstruction, taking out its signal during testing resulted in increase of the loss compared to the baseline. We z-scored the resulting difference values and plotted them on the brain.

5.6. Audio contamination analysis

Recent work has demonstrated that direct neural recordings can show signs of acoustic contamination in their acquired signals [57]. To account for this possibility, we performed an audio contamination analysis. Similar to the study by Roussel *et al* we computed correlations between spectrotemporal signals of microphone speech recordings and HD ECoG activity. Spectrograms were computed using the *librosa* (<https://librosa.github.io/librosa>) package. For this, short-time Fourier transform (STFT) was applied to audio (at sampling frequency of 16 kHz) and brain (at sampling frequency of 2 kHz) data. We selected signal processing parameters that produced spectrotemporal signals of matching temporal resolution (STFT window of 131 ms for audio and STFT window of 256 ms for HD ECoG), necessary for the correlation analysis. A default window overlap of 25% and a Hann window function were used. Prior to STFT calculation, HD ECoG data were notch-filtered to account for the line noise (at 50 Hz and all its harmonics) and referenced using either the common average or the bipolar scheme, since both types of signals had been used in the previous analyses. Pearson correlation coefficient was computed for audio and HD ECoG signals across all speech production trials (between 116 and 120 trials depending on the participant).

The analysis did not reveal any signs of acoustic contamination in four subjects out of five (figure 9). Only S5 showed noticeable correlations between several HD ECoG channels and the audio spectrogram in the range of 100–140 and 210–280 Hz. To perform a more detailed analysis we used the acoustic contamination toolbox provided by Roussel and colleagues [57]. Using the toolbox, we calculated cross-correlations between audio and HD ECoG signals of S5 during speech and calculated the statistical p -value for audio-ECoG correlations along the main diagonal. The detailed analysis revealed that (1) the bulk of the correlation peaked at the negative lag of 170 ms (ECoG activity preceded audio), and (2) the mean diagonal correlation of .2 did not reach significance based on the permutation tests (with the statistical p -value of $p = .19$, based on 10 000 reshuffles of contamination matrices).

5.7. Code and data availability

Code of all analyses can be accessed at https://github.com/Immiora/word_decoding_HD_ECoG.




Figure 9. Results of the acoustic contamination analysis. a. Spectrograms of the microphone audio recordings and associated HD ECoG signals. Per subject ECoG signal of one electrode is shown. The electrode was selected based on t -tests that compared mean HFB activity between speech and silence trials per electrode. The electrode with the largest positive t -value was used for the plot. Selected parameters resulted in the matched temporal resolution of the spectrogram signals (16000/(4096/4) in audio and 2000/(512/4) in HD ECoG). The audio spectrogram is shown in the frequency range that matches the ECoG spectrogram. Plots show HD ECoG data prior to referencing (common average or bipolar). (b), (c) Correlation plots between audio and HD ECoG temporal signal per frequency bin (between 0 and 300 Hz) and per electrode. Each value is the average Pearson correlation over all word production trials. Correlation plots are shown for CAR (b) and bipolar (c) reference schemes for HD ECoG signals. The bottom right panel shows results of the detailed contamination analysis using a toolbox from Roussel and colleagues [57]. The panel shows results of the cross-correlation analyses that indicate that the bulk of the correlation was not centered around zero (when ECoG and audio signals are aligned in time), but was shifted toward the lag of -170 ms. Negative shifts mean that ECoG activity preceded audio data it was correlated with, which agrees with the previous reports of pre-activation of motor and premotor cortices prior to speaking. The formal analysis revealed that mean correlation between audio and ECoG data was .2, which did not reach significance based on the permutation tests (with the statistical p -value of $p = .19$, based on 10 000 reshuffles of the correlation matrices).

Data availability statement

Data from human perceptual judgments can be found at <https://doi.org/10.34973/c07k-v019>. HD ECoG data will be released pending a temporary embargo and the required de-identification procedures. Behavioural data that support the findings of this study are openly available [83].

Acknowledgments

This work was supported by the European Research Council (Advanced iConnect Project Grant ADV 320708), the Netherlands Organisation for Scientific Research (Language in Interaction Consortium, Gravitation Grant No. 024.001.006; and INTENSE Consortium, Grant No. 1761) and National Institute On Deafness and Other Communication Disorders of the National Institutes of Health (Grant No. U01DC016686). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We thank Frans Leijten, Cyrille Ferrier, Geert-Jan Huiskamp, Sandra van der Salm and Tineke Gebbink for help with collecting ECoG data; Peter Gosselaar and Peter van Rijen for implanting the electrodes; ECoG patients and participants of the behavioral experiments for their time and effort; Zuzanna Dedyk for help with implementing behavioral experiments; and the members of the UMC Utrecht BCI team for help with ECoG data collection.

Author contributions statement

All authors conceived the idea of the study and the approach and contributed to data acquisition, J B preprocessed the data, implemented the modeling approach, performed the analyses and wrote the initial draft of the paper. All authors reviewed the manuscript.

ORCID iDs

Julia Berezutskaya  <https://orcid.org/0000-0002-2005-8758>

Zachary V Freudenburg  <https://orcid.org/0000-0002-2790-0020>

Mariska J Vansteensel  <https://orcid.org/0000-0002-9252-5116>

Erik J Aarnoutse  <https://orcid.org/0000-0001-7648-250X>

Nick F Ramsey  <https://orcid.org/0000-0002-7136-259X>

Marcel A J van Gerven  <https://orcid.org/0000-0002-2206-9098>

References

- [1] Blakely T, Miller K J, Rao R P, Holmes M D and Ojemann J G 2008 Localization and classification of phonemes using high spatial resolution electrocorticography (ECoG) grids *2008 30th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (IEEE)* pp 4964–7
- [2] Kellis S *et al* 2010 Decoding spoken words using local field potentials recorded from the cortical surface *J. Neural Eng.* **7** 056007
- [3] Leuthardt E C *et al* 2011 Using the electrocorticographic speech network to control a brain–computer interface in humans *J. Neural Eng.* **8** 036004
- [4] Pei X, Barbour D L, Leuthardt E C and Schalk G 2011 Decoding vowels and consonants in spoken and imagined words using electrocorticographic signals in humans *J. Neural Eng.* **8** 046028
- [5] Wang W, Degenhart A D, Sudre G P, Pomerleau D A and Tyler-Kabara E C 2011 Decoding semantic information from human electrocorticographic (ECoG) signals *2011 Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (IEEE)* pp 6294–8
- [6] Tankus A, Fried I and Shoham S 2012 Structured neuronal encoding and decoding of human speech features *Nat. Commun.* **3** 1–5
- [7] Pasley B N *et al* 2012 Reconstructing speech from human auditory cortex *PLoS Biol.* **10** e1001251
- [8] Ikeda S *et al* 2014 Neural decoding of single vowels during covert articulation using electrocorticography *Front. Hum. Neurosci.* **8** 125
- [9] Bouchard K E and Chang E F 2014 Neural decoding of spoken vowels from human sensory-motor cortex with high-density electrocorticography *2014 36th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (IEEE)* pp 6782–5
- [10] Mugler E M *et al* 2014 Direct classification of all american english phonemes using signals from functional speech motor cortex *J. Neural Eng.* **11** 035015
- [11] Martin S *et al* 2014 Decoding spectrotemporal features of overt and covert speech from the human cortex *Front. Neuroeng.* **7** 14
- [12] Herff C *et al* 2015 Brain-to-text: decoding spoken phrases from phone representations in the brain *Front. Neurosci.* **9** 217
- [13] Lotte F *et al* 2015 Electrocorticographic representations of segmental features in continuous speech *Front. Hum. Neurosci.* **9** 97
- [14] Martin S *et al* 2016 Word pair classification during imagined speech using direct brain recordings *Sci. Rep.* **6** 1–12
- [15] Ramsey N F *et al* 2018 Decoding spoken phonemes from sensorimotor cortex with high-density ECoG grids *Neuroimage* **180** 301–11
- [16] Ibayashi K *et al* 2018 Decoding speech with integrated hybrid signals recorded from the human ventral motor cortex *Front. Neurosci.* **12** 221
- [17] Livezey J A, Bouchard K E and Chang E F 2019 Deep learning as a tool for neural data analysis: speech classification and cross-frequency coupling in human sensorimotor cortex *PLoS Comput. Biol.* **15** e1007091
- [18] Anumanchipalli G K, Chartier J and Chang E F 2019 Speech synthesis from neural decoding of spoken sentences *Nature* **568** 493–8
- [19] Akbari H, Khalighinejad B, Herrero J L, Mehta A D and Mesgarani N 2019 Towards reconstructing intelligible speech from the human auditory cortex *Sci. Rep.* **9** 1–12
- [20] Angrick M *et al* 2019 Speech synthesis from ECoG using densely connected 3D convolutional neural networks *J. Neural Eng.* **16** 036019

- [21] Herff C *et al* 2019 Generating natural, intelligible speech from brain activity in motor, premotor and inferior frontal cortices *Front. Neurosci.* **13** 1267
- [22] Moses D A, Leonard M K, Makin J G and Chang E F 2019 Real-time decoding of question-and-answer speech dialogue using human cortical activity *Nat. Commun.* **10** 1–14
- [23] Makin J G, Moses D A and Chang E F 2020 Machine translation of cortical activity to text with an encoder–decoder framework *Nat. Neurosci.* **23** 575–82
- [24] Sun P, Anumanchipalli G K and Chang E F 2020 Brain2char: a deep architecture for decoding text from brain recordings *J. Neural Eng.* **17** 066015
- [25] Wilson G H *et al* 2020 Decoding spoken english from intracortical electrode arrays in dorsal precentral gyrus *J. Neural Eng.* **17** 066007
- [26] Kohler J *et al* 2021 Synthesizing speech from intracranial depth electrodes using an encoder-decoder framework (arXiv:2111.01457)
- [27] Angrick M *et al* 2021 Real-time synthesis of imagined speech processes from minimally invasive recordings of neural activity *Commun. Biol.* **4** 1–10
- [28] Proix T *et al* 2022 Imagined speech can be decoded from low-and cross-frequency intracranial EEG features *Nat. Commun.* **13** 1–14
- [29] Guenther F H *et al* 2009 A wireless brain-machine interface for real-time speech synthesis *PLoS One* **4** e8218
- [30] Brumberg J S, Wright E J, Andreassen D S, Guenther F H and Kennedy P R 2011 Classification of intended phoneme production from chronic intracortical microelectrode recordings in speech motor cortex *Front. Neurosci.* **5** 65
- [31] Vansteensel M J *et al* 2016 Fully implanted brain–computer interface in a locked-in patient with als *New Engl. J. Med.* **375** 2060–6
- [32] Stavisky S D *et al* 2019 Neural ensemble dynamics in dorsal motor cortex during speech in people with paralysis *eLife* **8** e46015
- [33] Moses D A *et al* 2021 Neuroprosthesis for decoding speech in a paralyzed person with anarthria *New Engl. J. Med.* **385** 217–27
- [34] Wandelt S K *et al* 2022 Decoding grasp and speech signals from the cortical grasp circuit in a tetraplegic human *Neuron* **110** 1777–87
- [35] Metzger S L *et al* 2022 Generalizable spelling using a speech neuroprosthesis in an individual with severe limb and vocal paralysis *Nat. Commun.* **13** 6510
- [36] Brumberg J S, Nieto-Castanon A, Kennedy P R and Guenther F H 2010 Brain–computer interfaces for speech communication *Speech Commun.* **52** 367–79
- [37] Chaudhary U, Birbaumer N and Ramos-Murguialday A 2016 Brain–computer interfaces for communication and rehabilitation *Nat. Rev. Neurol.* **12** 513–25
- [38] Bocquelet F, Hueber T, Girin L, Chabardès S and Yvert B 2016 Key considerations in designing a speech brain-computer interface *J. Physiol. Paris* **110** 392–401
- [39] Herff C and Schultz T 2016 Automatic speech recognition from neural signals: a focused review *Front. Neurosci.* **10** 429
- [40] Iljina O *et al* 2017 Neurolinguistic and machine-learning perspectives on direct speech BCIs for restoration of naturalistic communication *Brain-Comput. Interfaces* **4** 186–99
- [41] Cooney C, Folli R and Coyle D 2018 Neurolinguistics research advancing development of a direct-speech brain-computer interface *iScience* **8** 103–25
- [42] Martin S, Iturrate I, Millán J d R, Knight R T and Pasley B N 2018 Decoding inner speech using electrocorticography: progress and challenges toward a speech prosthesis *Front. Neurosci.* **12** 422
- [43] Martin S, Millán J d R, Knight R T and Pasley B N 2019 The use of intracranial recordings to decode human language: challenges and opportunities *Brain Lang.* **193** 73–83
- [44] Rabbani Q, Milsap G and Crone N E 2019 The potential for a speech brain–computer interface using chronic electrocorticography *Neurotherapeutics* **16** 144–65
- [45] Herff C, Krusienski D J and Kubben P 2020 The potential of stereotactic-EEG for brain-computer interfaces: current progress and future directions *Front. Neurosci.* **14** 123
- [46] Berezutskaya J, Saive A-L, Jerbi K and van Gerven M A J 2022 How does artificial intelligence contribute to iEEG research? (arXiv:2207.13190)
- [47] Bruurmijn M L, Pereboom I P, Vansteensel M J, Raemaekers M A and Ramsey N F 2017 Preservation of hand movement representation in the sensorimotor areas of amputees *Brain* **140** 3166–78
- [48] Collinger J L *et al* 2013 High-performance neuroprosthetic control by an individual with tetraplegia *Lancet* **381** 557–64
- [49] Chartier J, Anumanchipalli G K, Johnson K and Chang E F 2018 Encoding of articulatory kinematic trajectories in human speech sensorimotor cortex *Neuron* **98** 1042–54
- [50] Favero P, Berezutskaya J, Ramsey N F, Nazarov A and Freudenburg Z V 2022 Mapping acoustics to articulatory gestures in Dutch: relating speech gestures, acoustics and neural data 2022 44th Annual Int. Conf. IEEE Engineering in Medicine & Biology Society (EMBC) (IEEE) pp 802–6
- [51] Bocquelet F, Hueber T, Girin L, Savariaux C and Yvert B 2016 Real-time control of an articulatory-based speech synthesizer for brain computer interfaces *PLoS Comput. Biol.* **12** e1005119
- [52] Klein E, Kinsella M, Stevens I and Fried-Oken M 2022 Ethical issues raised by incorporating personalized language models into brain-computer interface communication technologies: a qualitative study of individuals with neurological disease *Disabil. Rehabil.: Assist. Technol.* **1–11**
- [53] Cheung C, Hamilton L S, Johnson K and Chang E F 2016 The auditory representation of speech sounds in human motor cortex *eLife* **5** e12577
- [54] Berezutskaya J, Baratin C, Freudenburg Z V and Ramsey N F 2020 High-density intracranial recordings reveal a distinct site in anterior dorsal precentral cortex that tracks perceived speech *Hum. Brain Mapp.* **41** 4587–609
- [55] Willett F R, Avansino D T, Hochberg L R, Henderson J M and Shenoy K V 2021 High-performance brain-to-text communication via handwriting *Nature* **593** 249–54
- [56] Livezey J A and Glaser J I 2021 Deep learning approaches for neural decoding across architectures and recording modalities *Briefings Bioinform.* **22** 1577–91
- [57] Roussel P *et al* 2020 Observation and assessment of acoustic contamination of electrophysiological brain signals during speech production and sound perception *J. Neural Eng.* **17** 056028
- [58] Huang G, Liu Z, Van Der Maaten L and Weinberger K Q 2017 Densely connected convolutional networks *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* pp 4700–8
- [59] Sutskever I, Vinyals O and Le Q V 2014 Sequence to sequence learning with neural networks *Advances in Neural Information Processing Systems* vol 27
- [60] Akiba T, Sano S, Yanase T, Ohta T and Koyama M 2019 Optuna: a next-generation hyperparameter optimization framework *Proc. 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* pp 2623–31
- [61] Stevens S S, Volkman J and Newman E B 1937 A scale for the measurement of the psychological magnitude pitch *J. Acoust. Soc. Am.* **8** 185–90
- [62] Taal C H, Hendriks R C, Heusdens R and Jensen J 2010 A short-time objective intelligibility measure for time-frequency weighted noisy speech 2010 IEEE International Conf. on Acoustics, Speech and Signal Processing (IEEE) pp 4214–7
- [63] Wairagkar M, Hochberg L R, Brandman D M and Stavisky S D 2023 Synthesizing speech by decoding intracortical neural activity from dorsal motor cortex 2023 11th Int. IEEE/EMBS Conf. on Neural Engineering (NER) (IEEE) pp 1–4

- [64] Yamamoto R, Song E and Kim J-M 2020 Parallel WaveGAN: a fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram *ICASSP 2020-2020 IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE) pp 6199–203
- [65] Spoerer C J, McClure P and Kriegeskorte N 2017 Recurrent convolutional neural networks: a better model of biological object recognition *Front. Psychol.* **8** 1551
- [66] Kietzmann T C et al 2019 Recurrence is required to capture the representational dynamics of the human visual system *Proc. Natl Acad. Sci.* **116** 21854–63
- [67] Berezutskaya J, Freudenburg Z V, Güçlü U, van Gerven M A J and Ramsey N F 2020 Brain-optimized extraction of complex sound features that drive continuous auditory perception *PLoS Comput. Biol.* **16** e1007992
- [68] Elango V, Patel A N, Miller K J and Gilja V 2017 Sequence transfer learning for neural decoding *bioRxiv* <https://doi.org/10.1101/210732> (posted online 23 December 2017)
- [69] Peterson S M, Steine-Hanson Z, Davis N, Rao R P and Brunton B W 2021 Generalized neural decoders for transfer learning across participants and recording modalities *J. Neural Eng.* **18** 026014
- [70] Wang R et al 2020 Stimulus speech decoding from human cortex with generative adversarial network transfer learning *2020 IEEE 17th Int. Symp. on Biomedical Imaging (ISBI)* (IEEE) pp 390–4
- [71] Berezutskaya J, Ambrogioni L, Ramsey N F and van Gerven M A 2022 Towards naturalistic speech decoding from intracranial brain data *2022 44th Annual Int. Conf. IEEE Engineering in Medicine & Biology Society (EMBC)* (IEEE) pp 3100–4
- [72] Nourski K V et al 2015 Sound identification in human auditory cortex: differential contribution of local field potentials and high gamma power as revealed by direct intracranial recordings *Brain Lang.* **148** 37–50
- [73] Hermes D, Miller K J, Noordmans H J, Vansteensel M J and Ramsey N F 2010 Automated electrocorticographic electrode localization on individually rendered brain surfaces *J. Neurosci. Methods* **185** 293–8
- [74] Branco M P et al 2018 Alice: a tool for automatic localization of intra-cranial electrodes for clinical and high-density grids *J. Neurosci. Methods* **301** 43–51
- [75] Fischl B 2012 FreeSurfer *Neuroimage* **62** 774–81
- [76] Boersma P 2001 Praat, a system for doing phonetics by computer *Glott. Int.* **5** 341–5
- [77] Luong M-T, Pham H and Manning C D 2015 Effective approaches to attention-based neural machine translation (arXiv:1508.04025)
- [78] Cho K, Van Merriënboer B, Bahdanau D and Bengio Y 2014 On the properties of neural machine translation: encoder-decoder approaches (arXiv:1409.1259)
- [79] Chung J, Gulcehre C, Cho K and Bengio Y 2014 Empirical evaluation of gated recurrent neural networks on sequence modeling (arXiv:1412.3555)
- [80] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)
- [81] Srivastava N, Hinton G, Krizhevsky A, Sutskever I and Salakhutdinov R 2014 Dropout: a simple way to prevent neural networks from overfitting *J. Mach. Learn. Res.* **15** 1929–58
- [82] Hutter F, Hoos H and Leyton-Brown K 2014 An efficient approach for assessing hyperparameter importance *Int. Conf. on Machine Learning* (PMLR) pp 754–62
- [83] Berezutskaya J and van Gerven M A J 2023 Behavioral assessment of the quality of speech reconstructions from intracranial neural activity Version 1 (Radboud University) (dataset) (available at: <https://doi.org/10.34973/c07k-v019>)