

Digital Voicing of Silent Speech

David Gaddy and Dan Klein

University of California, Berkeley

{dgaddy, klein}@berkeley.edu

Abstract

In this paper, we consider the task of digitally voicing silent speech, where silently mouthed words are converted to audible speech based on electromyography (EMG) sensor measurements that capture muscle impulses. While prior work has focused on training speech synthesis models from EMG collected during *vocalized* speech, we are the first to train from EMG collected during silently articulated speech. We introduce a method of training on silent EMG by transferring audio targets from vocalized to silent signals. Our method greatly improves intelligibility of audio generated from silent EMG compared to a baseline that only trains with vocalized data, decreasing transcription word error rate from 64% to 4% in one data condition and 88% to 68% in another. To spur further development on this task, we share our new dataset of silent and vocalized facial EMG measurements.

1 Introduction

In this paper, we are interested in enabling speech-like communication without requiring sound to be produced. By using muscular sensor measurements of speech articulator movement, we aim to capture silent speech - utterances that have been articulated without producing sound. In particular, we focus on the task which we call *digital voicing*, or generating synthetic speech to be transmitted or played back.

Digitally voicing silent speech has a wide array of potential applications. For example, it could be used to create a device analogous to a Bluetooth headset that allows people to carry on phone conversations without disrupting those around them. Such a device could also be useful in settings where the environment is too loud to capture audible speech or where maintaining silence is important. Alternatively, the technology could be used by some




Figure 1: Electromyography (EMG) electrodes placed on the face can detect muscle movements from speech articulators.

people who are no longer able to produce audible speech, such as individuals whose larynx has been removed due to trauma or disease (Meltzner et al., 2017). In addition to these direct uses of digital voicing for silent speech, it may also be useful as a component technology for creating silent speech-to-text systems (Schultz and Wand, 2010), making silent speech accessible to our devices and digital assistants by leveraging existing high-quality audio-based speech-to-text systems.

To capture information about articulator movement, we make use of surface electromyography (EMG). Surface EMG uses electrodes placed on top of the skin to measure electrical potentials caused by nearby muscle activity. By placing electrodes around the face and neck, we are able to capture signals from muscles in the speech articulators. Figure 1 shows the EMG electrodes used to capture signals, and Figure 2 shows an example of EMG signals captured. We collect EMG measurements during both vocalized speech (normal speech production that has voicing, friction, and other speech

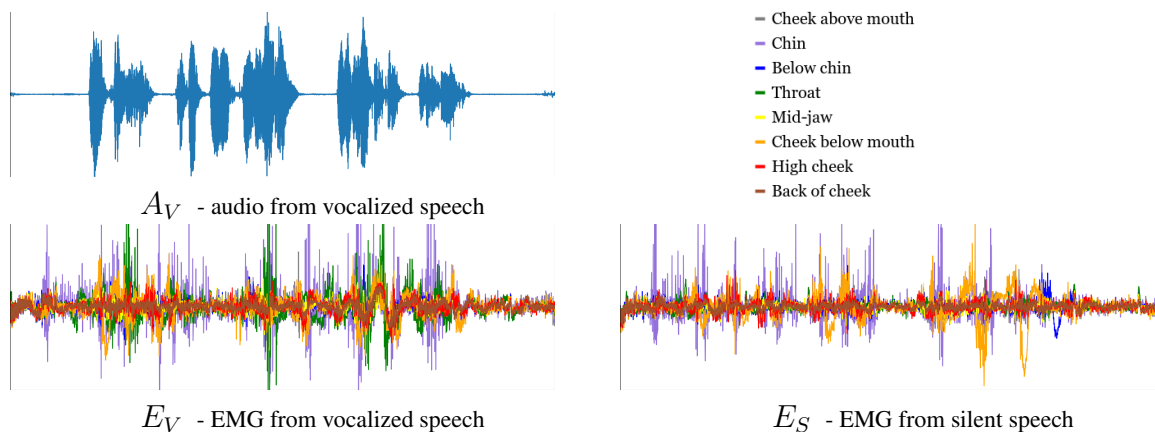


Figure 2: The three components of our data that we will use in our model. The vocalized speech signals A_V and E_V are collected simultaneously and so are time-aligned, while the silent signal E_S is a separate recording of the same utterance without vocalization. During training we use all three signals, and during testing we are given just E_S , from which we must generate audio. Colors represent different electrodes in the EMG data. Note that the silent EMG signal E_S is qualitatively different from its vocalized counterpart E_V . Not pictured, but also included in our data are the utterance texts, in this case: “It is possible that the infusoria under the microscope do the same.” (from H.G. Well’s *The War of the Worlds*).

sounds) and silent speech (speech-like articulations which do not produce sound). We denote these EMG signals E_V and E_S , respectively. During the vocalized speech we can also record audio A_V , but during silent speech there is no meaningful audio to record.

A substantial body of prior work has explored the use of facial EMG for silent speech-to-text interfaces (Jou et al., 2006; Schultz and Wand, 2010; Kapur et al., 2018; Meltzner et al., 2018). Several initial attempts have also been made to convert EMG signals to speech, similar to the task we approach in this paper (Toth et al., 2009; Janke and Diener, 2017; Diener et al., 2018). However, these works have focused on the artificial task of recovering audio from EMG that was recorded during *vocalized* speech, rather than the end-goal task of generating from silent speech. In terms of signals in Figure 2, prior work learned a model for producing audio A_V from vocalized EMG E_V and tested primarily on other vocalized EMG signals. While one might hope that a model trained in this way could directly transfer to silent EMG E_S , Toth et al. (2009) show that such a transfer causes a substantial degradation in quality, which we confirm in Section 4. This direct transfer from vocalized models fails to account for differences between features of the two speaking modes, such as a lack of voicing in the vocal folds and other changes in articulation to suppress sound.

In this paper, we extend digital voicing to train

on silent EMG E_S rather than only vocalized EMG E_V . Training with silent EMG is more challenging than with vocalized EMG, because when training on vocalized EMG data we have both EMG inputs and time-aligned speech targets, but for silent EMG any recorded audio will be silent. Our solution is to adopt a target-transfer approach, where audio output targets are transferred from vocalized recordings to silent recordings of the same utterances. We align the EMG features of the instance pairs with dynamic time warping (Rabiner and Juang, 1993), then make refinements to the alignments using canonical correlation analysis (Hotelling, 1936) and audio feature outputs from a partially trained model. The alignments can then be used to associate speech outputs with the silent EMG signals E_S , and these speech outputs are used as targets for training a recurrent neural transduction model.

We validate our method using both human and automatic metrics, and find that a model trained with our target transfer approach greatly outperforms a model trained on vocalized EMG alone. On a closed-vocabulary domain (date and time expressions §2.1), transcription word error rate (WER) from a human evaluation improves from 64% to just 4%. On a more challenging open vocabulary domain (reading from books §2.2) intelligibility measurements improve by 20% – from 88% to 68% with automatic transcription or 95% to 75% with human transcription.

We release our dataset of EMG signals collected

during both silent and vocalized speech. The dataset contains nearly 20 hours of facial EMG signals from a single speaker. To our knowledge, the largest public EMG-speech dataset previously available contains just 2 hours of data (Wand et al., 2014), and many papers continue to use private datasets. We hope that this public release will encourage development on the task and allow for fair comparisons between methods.

2 Data Collection

We collect a dataset of EMG signals and time-aligned audio from a single speaker during both silent and vocalized speech. Figure 2 shows an example from the data collected. The primary portion of the dataset consists of parallel silent / vocalized data, where the same utterances are recorded using both speaking modes. These examples can be viewed as tuples (E_S, E_V, A_V) of silent EMG, vocalized EMG, and vocalized audio, where E_V and A_V are time-aligned. Both speaking modes of an utterance were collected within a single session to ensure that electrode placement is consistent between them. For some utterances, we record only the vocalized speaking mode. We refer to these instances as non-parallel data, and represent them with the tuple (E_V, A_V) . Examples are segmented at the utterance level. The text that was read is included with each instance in the dataset, and is used as a reference when evaluating intelligibility in Section 4.

For comparison, we record data from two domains: a closed vocabulary and open vocabulary condition, which are described in Sections 2.1 and 2.2 below. Section 2.3 then provides additional details about the recording setup.

2.1 Closed Vocabulary Condition

Like other speech-related signals, the captured EMG signals from a particular phoneme may look different depending on its context. For this reason, our initial experiments will use a more focused vocabulary set before expanding to a large vocabulary in Section 2.2 below.

To create a closed-vocabulary data condition, we generate a set of date and time expressions for reading. These expressions come from a small set of templates such as “<weekday> <month> <year>” which are filled in with randomly selected values (over 50,000 unique utterances are possible from this scheme). Table 1 summarizes

Closed Vocabulary Condition

Parallel silent / vocalized speech

(E_S, E_V, A_V)

26 minutes silent / 30 minutes vocalized

Single session

500 utterances

Average of 4 words per utterance

67 words in vocabulary

Table 1: Closed vocabulary data summary

the properties of the data collected in this condition. A validation set of 30 utterances and a test set of 100 utterances are selected randomly, leaving 370 utterances for training.

2.2 Open Vocabulary Condition

The majority of our data was collected with open-vocabulary sentences from books. We use public domain books from Project Gutenberg.¹ Unlike the closed-vocabulary data which is collected in a single sitting, the open-vocabulary data is broken into multiple sessions where electrodes are reattached before each session and may have minor changes in position between different sessions. In addition to sessions with parallel silent and vocalized utterances, we also collect non-parallel sessions with only vocalized utterances. A summary of dataset features is shown in Table 2. We select a validation and test set randomly from the silent parallel EMG data, with 30 and 100 utterances respectively. Note that during testing, we use only the silent EMG recordings E_S , so the vocalized recordings of the test utterances are unused.

2.3 Recording Details

EMG signals are recorded using an OpenBCI Cyton Biosensing Board and transmitted to a computer over WiFi. Eight channels are collected at a sample rate of 1000 Hz. The electrode locations are described in Table 3. Gold-plated electrodes are used with Ten20 conductive electrode paste. We use a monopolar electrode configuration, with a shared reference electrode behind one ear. An electrode connected to the Cyton board’s bias pin is placed behind the other ear to actively cancel common-mode interference. A high pass Butterworth filter with cutoff 2 Hz is used to remove offset and drift in the collected signals, and AC

¹<https://www.gutenberg.org/>

Open Vocabulary Condition

Parallel Silent / Vocalized Speech

(E_S, E_V, A_V)

3.6 hours silent / 3.9 hours vocalized

Average session has 30 min. of each mode
1588 utterances

Non-parallel Vocalized Speech

(E_V, A_V)

11.2 hours

Average session length 67 minutes
5477 utterances

Total

18.6 hours

Average of 16 words per utterance
9828 words in vocabulary

Table 2: Open vocabulary data summary

	Location
1	left cheek just above mouth
2	left corner of chin
3	below chin back 3 cm
4	throat 3 cm left from Adam’s apple
5	mid-jaw right
6	right cheek just below mouth
7	right cheek 2 cm from nose
8	back of right cheek, 4 cm in front of ear
ref	below left ear
bias	below right ear

Table 3: Electrode locations.

electrical noise is removed with notch filters at 60 Hz and its harmonics. Forward-backward filters are used to avoid phase delay.

Audio is recorded from a built-in laptop microphone at 16kHz. Background noise is reduced using a spectral gating algorithm,² and volume is normalized across sessions based on peak root-mean-square levels.

3 Method

Our method is built around a recurrent neural transduction model from EMG features to time-aligned speech features (Section 3.1). We will denote the featurized version of the signals used by the transduction model $E'_{S/V}$ and A'_V for EMG and audio respectively. When training solely on vocal-

²<https://pypi.org/project/noisereducer/>

ized EMG data (E'_V to A'_V), training this model is straightforward. However, our experiments show that training on vocalized EMG alone leads to poor performance when testing on silent EMG (Section 4) because of differences between the two speaking modes.

A core contribution of our work is a method of training the transducer model on silent EMG signals, which no longer have time-aligned audio to use as training targets. We briefly describe our method here, then refer to section Section 3.2 for more details. Using a set of utterances recorded in both silent and vocalized speaking modes, we find alignments between the two recordings and use them to associate speech features from the vocalized instance (A'_V) with the silent EMG E'_S . The alignment is initially found using dynamic time warping between EMG signals and then is refined using canonical correlation analysis (CCA) and predicted audio from a partially trained model.

Finally, to generate audio from predicted speech features, we use a WaveNet decoder, as described in Section 3.3.

3.1 EMG to Speech Feature Transducer

When converting EMG input signals to audio outputs, our first step is to use a bidirectional LSTM to convert between featurized versions of the signals, E' and A' . Both feature representations operate at the same frequency, 100 Hz, so that each EMG input $E'[i]$ corresponds to a single time-aligned output $A'[i]$. Our primary features for representing EMG signals are the time domain features from Jou et al. (2006), which are commonly used in the EMG-speech-to-text literature. After splitting the signal from each channel into low and high-frequency components (x_{low} and x_{high}) using a triangular filter with cutoff 134 Hz, the signal is windowed with a frame length of 27 ms and shift of 10 ms. For each frame, five features are computed as follows:

$$\left[\frac{1}{n} \sum_i (x_{low}[i])^2, \frac{1}{n} \sum_i x_{low}[i], \frac{1}{n} \sum_i (x_{high}[i])^2, \frac{1}{n} \sum_i |x_{high}[i]|, \text{ZCR}(x_{high}) \right]$$

where ZCR is the zero-crossing rate. In addition to the time domain features, we also append magnitude values from a 16-point Short-time Fourier transform for each 27 ms frame, which gives us

9 additional features. The two representations result in a total of 112 features to represent the 8 EMG channels. Speech is represented with 26 Mel-frequency cepstral coefficients (MFCCs) from 27 ms frames with 10 ms stride. All EMG and audio features are normalized to approximately zero mean and unit variance before processing. To help the model to deal with minor differences in electrode placement across sessions, we represent each session with a 32 dimensional session embedding and append the session embedding to the EMG features across all timesteps of an example before feeding into the LSTM.

The LSTM model itself consists of 3 bidirectional LSTM layers with 1024 hidden units, followed by a linear projection to the speech feature dimension. Dropout 0.5 is used between all layers, as well as before the first LSTM and after the last LSTM. The model is trained with a mean squared error loss against time-aligned speech features using the Adam optimizer. The initial learning rate is set to .001, and is decayed by half after every 5 epochs with no improvement in validation loss. We evaluate a loss on the validation set at the end of every epoch, and select the parameters from the epoch with the best validation loss as the final model.

3.2 Audio Target Transfer

To train the EMG to speech feature transducer, we need speech features that are time-aligned with the EMG features to use as target outputs. However, when training with EMG from silent speech, simultaneously-collected audio recordings do not have any audible speech to use as targets. In this section, we describe how parallel utterances, as described in Section 2, can be used to transfer audio feature labels from a vocalized recording to a silent one. More concretely, given a tuple (E'_S, E'_V, A'_V) of features from silent speech EMG, vocalized speech EMG, and vocalized speech audio, where E'_V and A'_V are collected simultaneously, we estimate a set of audio features \tilde{A}'_S that time-align with E'_S and represent the output that we would like our transduction network to predict. A diagram of the method can be found in Figure 3.

Our alignment will make use of dynamic time warping (DTW) (Rabiner and Juang, 1993), a dynamic programming algorithm for finding a minimum-cost monotonic alignment between two sequences s_1 and s_2 . DTW builds a table $d[i, j]$ of

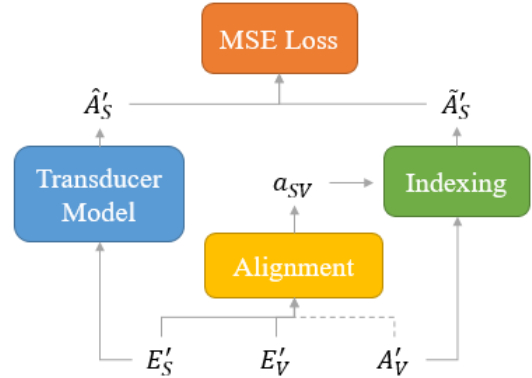


Figure 3: Our audio target transfer method for training on silent EMG E_S . Details in Section 3.2.

the minimum cost of alignment between the first i items in s_1 and the first j items in s_2 . The recursive step used to fill this table is $d[i, j] = \delta[i, j] + \min(d[i-1, j], d[i, j-1], d[i-1, j-1])$, where $\delta[i, j]$ is the local cost of aligning $s_1[i]$ with $s_2[j]$. After the dynamic program, we can follow backpointers through the table to find a path of (i, j) pairs representing an alignment. Although the path is monotonic, a single position i may repeat several times with increasing values of j . We take the first pair from any such sequence to form a mapping $a_{s_1 s_2}[i] \rightarrow j$ from every position i in s_1 to a position j in s_2 .

For our audio target transfer, we perform DTW as described above with $s_1 = E'_S$ and $s_2 = E'_V$. Initially, we use euclidean distance between the features of E'_S and E'_V for the alignment cost ($\delta_{\text{EMG}}[i, j] = \|E'_S[i] - E'_V[j]\|$), but will describe several refinements to this choice in Sections 3.2.1 and 3.2.2 below. DTW results in an alignment $a_{SV}[i] \rightarrow j$ that tells us a position j in E'_V for every position i in E'_S . We can then create a warped audio feature sequence \tilde{A}'_S that aligns with E'_S using $\tilde{A}'_S[i] = A'_V[a_{SV}[i]]$. During training of the EMG to audio transduction model, we will use \tilde{A}'_S as our targets for the transduction outputs \hat{A}'_S when calculating a loss.

This procedure of aligning signals to translate between them is reminiscent of some DTW-based methods for the related task of voice conversion (Kobayashi and Toda, 2018; Desai et al., 2009). The difference between these tasks is that our task operates on triples (E_S, E_V, A_V) and must account for the difference in modality between the input E_S and output A_V , while voice conversion operates in a single modality with examples of the form (A_1, A_2) .

In addition to training the transducer from E'_S to \hat{A}'_S , we also find that training on the vocalized signals (E'_V to A'_V) improves performance. The vocalized samples are labeled with different session embeddings to allow the model to specialize to each speaking mode. Each training batch contains samples from both modes mixed together. For the open vocabulary setting, the full set of examples to sample from has 3 sources: (E'_S, \hat{A}'_S) created from parallel utterances, (E'_V, A'_V) from the vocalized recording of the parallel utterances, and (E'_V, A'_V) from the non-parallel vocalized recordings.

3.2.1 CCA

While directly aligning EMG features E'_S and E'_V can give us a rough alignment between the signals, doing so ignores the differences between the two signals that lead us to want to train on the silent signals in the first place (e.g. inactivation of the vocal folds and changes in manner of articulation to prevent frication). To better capture correspondences between the signals, we use canonical correlation analysis (CCA) (Hotelling, 1936) to find components of the two signals which are more highly correlated. Given a number of paired vectors (v_1, v_2) , CCA finds linear projections P_1 and P_2 that maximize correlation between corresponding dimensions of $P_1 v_1$ and $P_2 v_2$.

To get the initial pairings required by CCA, we use alignments found by DTW with the raw EMG feature distance δ_{EMG} . We aggregate aligned E'_S and E'_V features over the entire dataset and feed these to a CCA algorithm to get projections P_S and P_V . CCA allows us to choose the dimensionality of the space we are projecting to, and we use 15 dimensions for all experiments. Using the projections from CCA, we define a new cost for DTW

$$\delta_{\text{CCA}}[i, j] = \left\| P_S E'_S[i] - P_V E'_V[j] \right\|$$

Our use of CCA for DTW is similar to Zhou and Torre (2009), which combined the two methods for use in aligning human pose data, but we found their iterative approach did not improve performance compared to a single application of CCA in our setting.

3.2.2 Refinement with Predicted Audio

So far, our alignments between the silent and vocalized recordings have relied solely on distances between EMG features. In this section, we propose an additional alignment distance term that uses audio features. Although the silent recording

has no useful audio signal, once we start to train a transducer model from E'_S to audio features, we can try to align the predicted audio features \hat{A}'_S to vocalized audio features A'_V . Combining with an EMG-based distance, our new cost for DTW becomes

$$\delta_{\text{full}}[i, j] = \delta_{\text{CCA}}[i, j] + \lambda \left\| \hat{A}'_S[i] - A'_V[j] \right\|$$

where λ is a hyperparameter to control the relative weight of the two terms. We use $\lambda = 10$ for all experiments in this paper.

When training a transducer model using predicted-audio alignment, we perform the first four epochs using only EMG-based alignment costs δ_{CCA} . Then, at the beginning of the fifth epoch, we use the partially-trained transducer model to compute alignments with cost δ_{full} . From then on, we re-compute alignments every five epochs of training.

3.3 WaveNet Synthesis

To synthesize audio from speech features, we use a WaveNet decoder (van den Oord et al., 2016), which generates the audio sample by sample conditioned on MFCC speech features A' . WaveNet is capable of generating fairly natural sounding speech, in contrast to the vocoder-based synthesizer used in previous EMG-to-speech papers, which caused significant degradation in naturalness (Janke and Diener, 2017). Our full synthesis model consists of a bidirectional LSTM of 512 dimensions, a linear projection down to 128 dimensions, and finally the WaveNet decoder which generates samples at 16 kHz. We use a WaveNet implementation from NVIDIA³ which provides efficient GPU inference. WaveNet hyperparameters can be found in Appendix A. During training, the model is given gold speech features as input, which we found to work better than training from EMG-predicted features. Due to memory constraints we do not use any batching during training, but other optimization hyperparameters are the same as those from Section 3.1.

4 Experiments

In this section, we run experiments to measure intelligibility of audio generated by our model from silent EMG signals E_S . Since prior work has trained only on vocalized EMG signals E_V , we

³<https://github.com/NVIDIA/nv-wavenet>

compare our method to a *direct transfer* baseline which trains a transducer model only on vocalized EMG E_V before testing on the silent EMG E_S .⁴ The baseline transducer and wavenet models have identical architecture to those used by our method, but are not trained with silent EMG using our target transfer approach. Since one may hypothesize that most of the differences between silent and vocalized EMG will take place near the vocal folds, we also test a variant of this baseline where the electrode placed on the neck is ignored.

We first test on the closed vocabulary data described in Section 2.1, then on the open vocabulary data from Section 2.2. On the open vocabulary data, we also run ablations to evaluate different alignment refinements with CCA and predicted audio (see Sections 3.2.1 and 3.2.2).

4.1 Closed Vocabulary Condition

We begin by testing intelligibility on the closed vocabulary data and time data with a human transcription evaluation. The human evaluator is given a set of 20 audio output files from each model being tested (listed below) and is asked to write out in words what they heard. The files to transcribe are randomly shuffled, and the evaluator is not told that the outputs come from different systems. They are told that the examples will contain dates and times, but are not given any further information about what types of expressions may occur. The full text of the instructions provided to the evaluator can be found in Appendix B. We compare the transcriptions from the human evaluator to the original text prompts that were read during data collection to compute a transcription word error rate (WER):

$$\text{WER} = \frac{\text{substitutions} + \text{insertions} + \text{deletions}}{\text{reference length}}$$

Lower WER values indicate better models.

Using this evaluation, we compare three different models: a direct transfer baseline trained only on vocalized EMG signals, a variant of this baseline where the throat electrode is removed to reduce divergence between speaking modes, and our full model trained on silent EMG using target-transfer. All three models were trained on open vocabulary

⁴Note that because prior work has not released data or code, we are unable to perform a direct comparison to experiments found in their papers. Our direct transfer baseline represents a conceptually equivalent model, but with larger neural networks than prior work.

Model	WER
Direct transfer baseline	88.8
Without throat electrode	64.6
Our model	3.6

Table 4: Results of a human intelligibility evaluation on the closed vocabulary data. Lower WER is better. Our model greatly outperforms both variants of the direct transfer baseline.

data (Section 2.2) before being fine-tuned on the closed vocabulary training set. A single WaveNet model is used to synthesize audio for all three models and was also trained on the open vocabulary data before being fine-tuned in-domain.

The results of our evaluation are shown in Table 4. We first observe that removing the throat electrode substantially improves intelligibility for the direct transfer baseline. Although this modification removes potentially useful information, it also removes divergence between the silent and vocalized EMG signals. Its relative success further motivates the need for methods to account for the differences in the two modes, such as our target-transfer approach. However, even with the throat-removal modification, the direct transfer approach is still only partially intelligible.

A model trained with our full approach, including CCA and predicted-audio alignment, achieves a WER of 3.6%. This result represents a high level of intelligibility and a 94% relative error reduction from the strongest baseline.

4.2 Open Vocabulary Condition

Similar to our evaluation in Section 4.1, we use a transcription WER to evaluate intelligibility of model outputs in the open vocabulary condition. For the open vocabulary setting, we evaluate both with a human transcription and with transcriptions from an automatic speech recognizer.

4.2.1 Human Evaluation

Our human evaluation with open vocabulary outputs follows the same setup as the closed vocabulary evaluation. Transcripts are collected for 20 audio outputs from each system, with a random interleaving of outputs from the different systems. The annotator had no prior information on the content of the texts being evaluated. We compare two systems: direct transfer without the throat electrode (the stronger baseline) and our full model.

Model	WER
Direct transfer baseline	91.2
Without throat electrode	88.0
Our model	68.0
Without CCA	69.8
Without audio alignment	76.5

Table 5: Results of an automatic intelligibility evaluation on open vocabulary data. Lower WER is better.

The results of this evaluation are a 95.1% WER for the direct transfer baseline and 74.8% WER for our system. While the intelligibility is much lower than in the closed vocabulary condition, our method still strongly out-performs the baseline with a 20% absolute improvement.

4.2.2 Automatic Evaluation

In addition to the human evaluation, we also perform an automatic evaluation by transcribing system outputs with a large-vocabulary automatic speech recognition (ASR) system. Using an automatic transcription allows for much faster and more reproducible comparisons between methods compared to a human evaluation. For our automatic speech recognizer, we use the open source implementation of DeepSpeech from Mozilla⁵ (Hannun et al., 2014). Running the recognizer on the original vocalized audio recordings from the test set results in a WER of 9.5%, which represents a lower bound for this evaluation.

Our automatic evaluation results are shown in Table 5. While the absolute WER values for the ASR evaluation do not perfectly match those of the human transcriptions, both evaluations show a 20% improvement of our system over the best baseline. Given this correlation between evaluations and the many advantages of automated evaluation, we will use the automatic metric throughout the rest of this work and recommend its use for comparisons in future work.

We also run ablations of the two alignment refinement methods from Sections 3.2.1 and 3.2.2 and include results in Table 5. We see that both refinements have a positive effect on performance, though the impact of aligning with predicted audio is greater.




Figure 4: Effect of data amount on intelligibility.

4.3 Additional Experiments

In the following subsections, we perform additional experiments on the open vocabulary data to explore the effect of data size and choice of electrode positions. These experiments are all evaluated using the automatic transcription method described in Section 4.2.

4.3.1 Data Size

In this section we explore the effect of dataset size on model performance. We train the EMG-to-speech transducer model on various-sized fractions of the dataset, from 10% to 100%, and plot the resulting WER. We select from the parallel (silent and vocalized) and non-parallel (vocalized only) portions proportionally here, but will re-visit the difference later. Although data size also affects WaveNet quality, we use a single WaveNet trained on the full dataset for all evaluations to focus on EMG-specific data needs.

Figure 4 shows the resulting intelligibility measurements for each data size. As would be expected, the rate of improvement is larger when data sizes are small. However, there does not seem to be a plateau in performance, as improvements continue even when increasing data size beyond fifteen hours. These continued gains suggest that collecting additional data could provide more improvement in the future.

We also train a model without the non-parallel vocalized data (vocalized recordings with no associated silent recording; see Section 2). A model trained without this data has a WER of 71.6%, a loss of 3.6 absolute percentage points. This confirms that non-parallel vocalized data can be useful for silent speech even though it contains only data from the vocalized speaking mode. However, if we compare this accuracy to a model where the same amount of data was removed proportionally from

⁵<https://github.com/mozilla/DeepSpeech>

the two data types (parallel and non-parallel), we see that removing a mixture of both types leads to a much larger performance decrease to 76% WER. This indicates that the non-parallel data is less important to the performance of our model, and suggests that future data collection efforts should focus on collecting parallel utterances of silent and vocalized speech rather than non-parallel utterances of vocalized speech.

4.3.2 Removing Electrodes

In this section, we experiment with models that operate on a reduced set of electrodes to assess the impact on performance and gain information about which electrodes are most important. We perform a random search to try to find a subset of four electrodes that works well. More specifically, we sample 10 random combinations of four electrodes to remove (out of 70 possible combinations) and train a model with each. We then use validation loss to select the best models.

The three best-performing models removed the following sets of electrodes (using electrode numbering from Table 3): 1) {4, 5, 7, 8} 2) {3, 5, 7, 8} and 3) {2, 5, 7, 8}. We note that electrodes 5, 7, and 8 (which correspond with electrodes on the mid-jaw, upper cheek, and back cheek) appear in all of these, indicating that they may be contributing less to the performance of the model. However, the best model we tested with four electrodes did have substantially worse intelligibility compared to an eight-electrode model, with 76.8% WER compared to 68.0%. A model that removed only electrodes 5, 7, and 8 also performed substantially worse, with a WER of 75.3%.

5 Conclusion

Our results show that digital voicing of silent speech, while still challenging in open domain settings, shows promise as an achievable technology. We show that it is important to account for differences in EMG signals between silent and vocalized speaking modes and demonstrate an effective method of doing so. On silent EMG recordings from closed vocabulary data our speech outputs achieve high intelligibility, with a 3.6% transcription word error rate and relative error reduction of 95% from our baseline. We also significantly improve intelligibility in an open vocabulary condition, with a relative error reduction over 20%. We hope that our public release of data will encourage

others to further improve models for this task.⁶

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 1618460.

References

- Srinivas Desai, E Veera Raghavendra, B Yegnanarayana, Alan W Black, and Kishore Prahallad. 2009. Voice conversion using artificial neural networks. In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3893–3896. IEEE.
- L. Diener, G. Felsch, M. Angrick, and T. Schultz. 2018. Session-independent array-based EMG-to-speech conversion using convolutional neural networks. In *Speech Communication; 13th ITG-Symposium*, pages 1–5.
- Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*.
- Harold Hotelling. 1936. Relations between two sets of variates.
- M. Janke and L. Diener. 2017. EMG-to-speech: Direct generation of speech from facial electromyographic signals. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(12):2375–2385.
- Szu-Chen Stan Jou, Tanja Schultz, Matthias Walliczek, Florian Kraft, and Alexander H. Waibel. 2006. Towards continuous speech recognition using surface electromyography. In *INTERSPEECH*.
- Arnav Kapur, Shreyas Kapur, and Pattie Maes. 2018. Alterego: A personalized wearable silent speech interface. In *23rd International Conference on Intelligent User Interfaces*, pages 43–53.
- Kazuhiro Kobayashi and Tomoki Toda. 2018. sprocket: Open-source voice conversion software. In *Odyssey*, pages 203–210.
- Geoffrey S Meltzner, James T Heaton, Yunbin Deng, Gianluca De Luca, Serge H Roy, and Joshua C Kline. 2017. Silent speech recognition as an alternative communication device for persons with laryngectomy. *IEEE/ACM transactions on audio, speech, and language processing*, 25(12):2386–2398.

⁶Our dataset can be downloaded from <https://doi.org/10.5281/zenodo.4064408> and code is available at https://github.com/dgaddy/silent_speech.

Geoffrey S Meltzner, James T Heaton, Yunbin Deng, Gianluca De Luca, Serge H Roy, and Joshua C Kline. 2018. Development of sEMG sensors and algorithms for silent speech recognition. *Journal of neural engineering*, 15(4):046031.

Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. 2016. WaveNet: A generative model for raw audio. *ArXiv*, abs/1609.03499.

Lawrence Rabiner and Biing-Hwang Juang. 1993. *Fundamentals of speech recognition*. Prentice Hall.

Tanja Schultz and Michael Wand. 2010. Modeling coarticulation in EMG-based continuous speech recognition. *Speech Communication*, 52(4):341–353.

Arthur R. Toth, Michael Wand, and Tanja Schultz. 2009. Synthesizing speech from electromyography using voice transformation techniques. In *INTER-SPEECH*.

Michael Wand, Matthias Janke, and Tanja Schultz. 2014. The EMG-UKA corpus for electromyographic speech processing. In *INTERSPEECH*.

Feng Zhou and Fernando Torre. 2009. Canonical time warping for alignment of human behavior. In *Advances in neural information processing systems*, pages 2286–2294.

A WaveNet Hyperparameters

Hyperparameter	Value
n_in_channels	256
n_layers	16
max_dilation	128
n_residual_channels	64
n_skip_channels	256
n_out_channels	256
n_cond_channels	128
upsamp_window	432
upsamp_stride	160

B Human Evaluator Instructions

The instructions given to the human evaluator are as follows: “Please listen to each of the attached sound files and write down what you hear as best you can. There are 60 files, each of which will contain an expression of some date or time. Write your transcriptions into a spreadsheet such as Excel or Google sheets so that the row numbers match the numbers in the file names. Although many of the clips will contain numbers, please write out what you hear as words. For example, you might write something like: five oh two pm

on Thursday⁷ Many of the clips may be difficult to hear. If this is the case, write whatever words you are able to make out, even if it does not form a complete expression. For example: five two pm on If you cannot make out any words, leave the corresponding row blank.”

C Additional Data Collection Details

During data collection, text prompts consisting of a single sentence to be read are displayed on a screen. After reading the sentence, the subject pressed a key to advance to the next sentence. If they were unhappy with a recording, they could press another key to re-record an utterance. A real-time display of EMG signals was used to monitor the electrodes for excessive noise. During silent speech, the subject was instructed to mouth words as naturally as possible without producing sound.

D Additional Reproducibility Information

Models were trained for up to two days on a single K80 GPU. Hyperparameter search consisted of a mixture of manual and random search, typically with less than 10 runs. Hyperparameters were chosen primarily based validation loss, with major design decisions also being checked with automatic transcription evaluation.

⁷We intentionally used an example that does not match a pattern in our generation procedure to avoid biasing the evaluator.