TAL: A SYNCHRONISED MULTI-SPEAKER CORPUS OF ULTRASOUND TONGUE IMAGING, AUDIO, AND LIP VIDEOS

Manuel Sam Ribeiro¹, Jennifer Sanger¹, Jing-Xuan Zhang^{1,2}, Aciel Eshky¹, Alan Wrench^{3,4}, Korin Richmond¹, Steve Renals¹

¹The Centre for Speech Technology Research, University of Edinburgh, UK, ²University of Science and Technology of China, P.R.China, ³Clinical Audiology, Speech and Language Research Centre, Queen Margaret University, UK, ⁴Articulate Instruments Ltd., UK

{sam.ribeiro, korin.richmond, s.renals}@ed.ac.uk

ABSTRACT

We present the Tongue and Lips corpus (TaL), a multi-speaker corpus of audio, ultrasound tongue imaging, and lip videos. TaL consists of two parts: TaL1 is a set of six recording sessions of one professional voice talent, a male native speaker of English; TaL80 is a set of recording sessions of 81 native speakers of English; taL80 out voice talent experience. Overall, the corpus contains 24 hours of parallel ultrasound, video, and audio data, of which approximately 13.5 hours are speech. This paper describes the corpus and presents benchmark results for the tasks of speech recognition, speech synthesis (articulatory-to-acoustic mapping), and automatic synchronisation of ultrasound to audio. The TaL corpus is publicly available under the CC BY-NC 4.0 license.

Index Terms— Ultrasound tongue imaging, video lip imaging, silent speech, articulography, corpora

1. INTRODUCTION

Measuring the position of the articulators during the speech production process is relevant to disciplines such as linguistics, speech processing, speech pathology, and anatomy. Articulatory movement can be captured [1] using *magnetic articulography techniques*, such as electromagnetic articulography (EMA) or permanent-magnetic articulography (PMA); *palatography techniques*, such as electropalatography or optopalatography; or *imaging techniques*, such as video imaging, magnetic resonance imaging (MRI) or ultrasound tongue imaging (UTI).

Video is the cheapest and most convenient imaging method to acquire articulatory data and is useful for studies focusing on the dynamics of those articulators. Its main limitation is its inability to capture anything beyond extraoral articulators, such as the lips and jaw. To provide complementary data, intraoral articulators can be monitored using medical imaging techniques such as MRI [2] and UTI [3]. Although MRI captures high-quality images, it suffers from a variety of disadvantages [4]: it is expensive and not easily accessible, and suffers from loud background noise, a supine recording position, and low temporal resolution. Ultrasound tongue imaging, on the other hand, is portable, non-invasive, clinically safe, and relatively cheap. UTI uses diagnostic ultrasound operating in B-mode to visualize the tongue surface during speech production at high frame rates. There are, however, some challenges associated with ultrasound images of the tongue [5]. They may be noisy and image resolution may not be ideal, as they are susceptible to unrelated highcontrast edges, speckle noise, or interruptions of the tongue surface. Furthermore, image quality may be affected by speaker characteristics such as age or physiology, or by probe placement.

Ultrasound tongue imaging has been used in various applications, including speech therapy [6, 7, 8], language learning [9, 10], phonetics studies [4], and the development of silent speech interfaces [11]. Previous work in the context of speech therapy used ultrasound data to develop tongue contour extractors [12], animate a tongue model [13], and automatically synchronise and process speech therapy recordings [14, 15]. Additionally, speech recognition [16, 17] and speech synthesis [18, 19] from ultrasound images have been used in silent speech interfaces to restore spoken communication for users with voice impairments or to allow silent communication in situations where audible speech is undesirable. Ultrasound data has also been used to develop models for articulatoryto-acoustic or acoustic-to-articulatory mapping [20, 21]. Lip videos have been used for research on multimodal speech perception [22], audio-visual synchronisation [23], automatic lip reading [24], audiovisual speech recognition [25], and speech reconstruction from silent videos [26].

There are relatively few publicly available corpora of ultrasound tongue images. The Silent Speech Challenge dataset [27] consists of approximately 2500 utterances of ultrasound and video images from a single native English speaker. The UltraSuite repository [28] contains ultrasound and audio data of 58 typically developing children and 28 children with speech sound disorders in the context of speech therapy. Because video imaging is easier to acquire, there is a much larger selection of audio-visual datasets of lip videos (see [24, 29] for recent surveys). Earlier corpora were concerned with restricted-vocabulary tasks, such as the recognition of the alphabet [30], isolated digits [31], or sentences with limited vocabulary [32]. Recent work, however, is based on more comprehensive audio-visual datasets containing several hundred hours of speech and thousands of speakers [33, 34].

This paper presents the Tongue and Lips Corpus (TaL), which contains synchronised imaging data of extraoral (lips) and intraoral

Supported by the Carnegie Trust for the Universities of Scotland (Research Incentive Grant number 008585) and the EPSRC Healthcare Partnerships grant number EP/P02338X/1 (Ultrax2020 — http://www.ultrax-speech.org).

Table 1. TaL corpus at a glance. Gender, accent, and age are self-reported. Accent is categorised here in terms of English (eng), Scottish (sco), or other.

	TaL1	TaL80
Number of speakers	1	81
Gender (male/female)	1/0	36/45
Accent (eng/sco/other)	1/0/0	40/36/5
Age (mean/std)	61/-	29.4/11.3
Utterances per session (mean/std)	263/14	205/25
Total number of utterances	1582	16639

(tongue) articulators from 82 native speakers of English. Section 2 describes the data collection and data post-processing steps of the TaL corpus. Section 3 introduces a set of benchmark experiments for automatic speech recognition, speech synthesis, and automatic synchronisation of ultrasound and audio. We conclude in Section 4 with an analysis of speaker performance across tasks and systems.

2. THE TAL CORPUS

The TaL corpus (Table 1) comprises **TaL1**, a set of six recording sessions of one professional voice talent, a male native speaker of English; and **TaL80**, a set of recording sessions of 81 native speakers of English, without voice talent experience. All recording sessions took place between October 2019 and March 2020. Sessions with the experienced voice talent were approximately 120 minutes in duration, and sessions with the remaining speakers were approximately 80 minutes.

Each participant was fitted with the UltraFit stabilising helmet [35], which held the video camera and the ultrasound probe. Data was recorded using the Articulate Assistant Advanced (AAA) software [36]. Ultrasound was recorded using Articulate Instruments' Micro system at ~80fps with a 92° field of view. A single B-Mode ultrasound frame has 842 echo returns for each of 64 scan lines, giving a 64×842 "raw" ultrasound frame that captures a midsagittal view of the tongue. Video images of the lips were recorded at ~60fps (greyscale interlaced) and synchronised with audio using Articulate Instruments' SynchBrightUp unit [37]. Participants were seated in a hemi-anechoic chamber and audio was captured with a Sennheiser HKH 800 p48 microphone with a 48KHz sampling frequency at 16 bit.

During data collection there was no attempt to normalise the position of the ultrasound probe or video camera across recording sessions. Instead, we attempted to position the ultrasound probe individually such that the best possible image was captured for each speaker. Similarly, the video camera was positioned such that the entirety of the lips was visible on screen. This is a similar process to what we might find in use cases of this technology, such as speech therapy [6]. There were, however, some challenges during data collection, as participants often moved when speaking or the stabilising helmet shifted. Whenever there was a substantial degradation in image quality, the equipment was manually adjusted, and the participants were asked to read a set of calibration utterances.

Recording prompts were extracted from a variety of sources: the accent identification paragraph [38], the Rainbow Passage [39], the Harvard sentences [40], the TIMIT corpus [41], the VCTK corpus [42], and the Librispeech corpus [43]. All prompts were spell-checked and adjusted for British English spelling. Table 2 shows a summary of the recording protocol used for each speaker in

Table 2. Recording protocol for TaL80. "Type" indicates speech type (read or silent speech). "Shared" indicates that the prompts were read by all speakers in TaL80. The last row does not have a defined number of prompts, as participants read as many as possible until the end of the session.

Prompts	Туре	Shared	Count
Swallow	-	Yes	1
Calibration	Read	Yes	2
Rainbow, Accent	Read	Yes	24
Calibration	Silent	Yes	2
Rainbow (partial)	Silent	Yes	6
Harvard	Silent	No	9
Harvard	Read	No	9
Spontaneous	-	No	1
TIMIT, VCTK, Librispeech	Read	No	-

TaL80. Participants were asked to swallow and to read two calibration sentences at the beginning and end of each recording session, and before and after any breaks. The protocol includes a mixture of prompts read by all speakers and prompts read by each speaker. Some prompts were read silently, audibly, or both. Additionally, each participant was asked to speak unprompted for 30-60 seconds. A list of suggested conversation topics was available, although participants were allowed to choose any topic of their preference. Each TaL1 session follows a similar protocol, with the addition of whispered speech utterances.

Using the AAA software, we synchronised the video and ultrasound data to the audio stream. During the synchronisation process, the video stream was deinterlaced. We then exported all recorded data from the AAA software. We normalised waveform levels using the tools available in [44]. Using *ffmpeg*, the video data was scaled to 240×320 pixels and encoded with *libx264* using *yuv420p* pixel format and a bit rate of 1 Mbps. Due to a configuration error, the synchronisation signal was corrupted for the first recording session of TaL1 (*day1*). For this reason, video synchronisation is not available for this session.

For each speaker, utterances were sorted by recording date-time and indexed from 001. Each filename contains an additional tag, denoting its prompt type. The tag **swa** is used for swallows, **cal** for calibration utterances, **spo** for spontaneous speech, **sil** for silent speech, and **aud** for audible read speech. The marker **x** was used with audible and silent tags to indicate cross-prompts (repeated by all speakers). The absence of this marker indicates that the prompt is unique to the current speaker. The TaL1 data contains the additional **whi** tag to indicate whispered speech. Cross-prompt tags in this data indicate prompts repeated across sessions. Furthermore, each utterance consists of five core data types identified by file extension:

- 1. The **prompt** file is identified by extension *.txt* and it includes the prompt text read by the speaker, and the date and time of the recording.
- 2. The **waveform**, identified by the file extension *.wav*, is a single-channel RIFF wave file, sampled at 48 KHz with a bit-depth of 16-bit.
- 3. The **synchronization signal** is stored identically to the speech waveform, but is identified by the file extension *.sync*. This waveform contains the audio pulses used to synchronise the video and the ultrasound stream.

Table 3. Amount of data for the TaL corpus by prompt type, computed over parallel data streams (audio, video, ultrasound). The *speech* columns estimate data after Voice Activity Detection. All estimates are given in minutes, except for the last row, which is given in hours. Shared data (denoted by x in the tag identifier) means that the same set of prompts are read by all speakers. For TaL1, x is applied across recording sessions.

TaL80	
n total	
3 829.45	
62.47	
-	
3 218.13	
78.19	
-	
56.24	
48.90	
hrs 21.90 hrs	

- 4. The **ultrasound** data is stored across two files. Raw ultrasound data is identified by the file extension *.ult*, while the extension *.param* is a text file containing ultrasound metadata (e.g. frames per second, synchronization offset, etc.).
- 5. The **video data** is identified by the file extension *.mp4*, which embeds its metadata.

We transcribed all spontaneous speech utterances and included the text transcription in the corresponding prompt file. Because spontaneous speech utterances can be long in duration (up to 60 seconds), we manually annotated the boundaries of shorter time segments (typically 5-10 seconds). This segmentation was added as an additional data type for each speaker, identified by the file extension *.lab*. This additional data type, available only for spontaneous speech utterances, is a text file with the start and end time in seconds for each time segment and their respective transcriptions.

The AAA software receives data from each modality separately. The ultrasound stream begins recording after the video and audio and normally stops recording after. We release all data streams without trimming them, as they may be useful for a variety of tasks which are not dependent on parallel data. Identifying parallel data should be trivial, given the available metadata for each data stream.

For a summary of the TaL corpus, we refer the reader back to Table 1. Table 3 includes more a detailed distribution of the amount of data by prompt. Overall, we have collected roughly 24 hours of parallel ultrasound, video, and audio, of which approximately 13.5 hours are speech. This was collected for a total of 82 speakers across 87 recording sessions. **TaL1** contains roughly 2 hours of data from a single experienced voice talent across six recording sessions. **TaL80** contains almost 22 hours of data from 81 native speakers of English. Figure 1 shows the distribution of ages for TaL80. There is a slight bias towards younger speakers, as participants were recruited primarily via University channels. Figure 2 shows sample frames from twelve speakers available in TaL80.

3. EXPERIMENTS

In this section, we present a set of benchmark experiments for the TaL corpus, focusing on the available modalities. Section 3.1 presents initial results for automatic speech recognition, Section 3.2 provides results for the articulatory-to-acoustic mapping task, and



Fig. 1. Age distribution for TaL80.

Section 3.3 presents results for the automatic synchronisation of ultrasound and audio.

3.1. Speech recognition

This set of experiments provides results for speech recognition using the three modalities available in the TaL corpus. We further consider these results across two training scenarios: speaker dependent and speaker independent. For this task, we use only read speech utterances (identified by *aud*), and discard silent, spontaneous, whispered, and other shared prompts. We expect speech styles to have different properties in terms of tongue and lip movement, therefore we focus only on read speech.

For the **speaker dependent** scenario, we conduct our experiments with the TaL1 dataset. We discard *day1*, for which video synchronisation is not available. The remaining sessions, days 2-6, are evaluated using cross-validation, leaving one recording session out for testing. We aggregate results by computing WER on the total number of errors and words rather than averaging WER across the 5 folds. For the **speaker independent** scenario, we train models on speakers 11-81 from the TaL80 dataset. We discard from the training set all prompts that co-occur with the TaL1 dataset. We save the first 10 speakers for testing, as these have the biggest prompt overlap with TaL1. Days 2-6 from TaL1 are used as test set, which is identical to the cross-validation setup for the speaker dependent system.

Feature extraction. During data preparation, calibration utterances are used to identify breaks in the recording session or adjustments to the equipment. This may indicate different image configurations, so we define a "speaker" to be any set of utterances without any breaks. This arrangement allow us to do speaker adaptive training on a speaker dependent scenario. We extract features over parallel data streams and we trim initial and end silences in the ultrasound and video modalities after Voice Activity Detection on the audio stream. For the audio stream, we extract 13 dimensional Mel Frequency Cepstral Coefficients (MFCCs) and we append their respective first and second derivatives. For the visual stream, we apply the 2D Discrete Cosine Transform (DCT) to each ultrasound or video frame. The upper left submatrix is flattened and used as input. Video frames are resized to 120×160 and ultrasound frames are resized to 64×420 . Both streams use a submatrix of 12×12 coefficients (144 coefficients).



Fig. 2. Sample frames from twelve speakers from the TaL corpus. The top figures show video images of the lips and the bottom figures show the corresponding ultrasound images of the tongue.

Table 4. Word Error Rate (WER) for speaker dependent and speaker independent speech recognition on TaL1 and TaL80.

Data	Speaker Dependent	Speaker Independent		
streams	TaL1	TaL1	TaL80	
Audio	3.42%	4.06%	13.00%	
Tongue	38.41%	75.34%	83.81%	
Lips	51.57%	84.70%	81.52%	
Tongue+Lips	25.09%	49.12%	59.20%	

Model training. We train the "acoustic" models with the Kaldi speech recognition toolkit [45]. For these experiments, we initialise models from a flat start using the corresponding features. After monophone and triphone training, input features are processed with Linear Discriminant Analysis (LDA) and a Maximum Likelihood Linear Transform (MLLT). This is followed by Speaker Adaptive Training (SAT) with feature-space MLLR (fMLLR) [46]. The alignment and HMM models from this stage are then used to train a time-delay neural network (TDNN, [47]) following Kaldi's *nnet3* recipe. We decode with a simple in-domain bigram language model trained on all prompts from the sources described in section 2.

Results and future work. Results are presented in Table 4. As expected, recognition results from the audio stream outperform those of the articulatory modalities by a large margin. Although we have made no attempts to design more complex systems for silent speech recognition, the high WERs illustrate the difficulty of the problem. This is particularly noticeable when using data from multiple speakers in a speaker independent system. In all cases, we observe that the visual data streams complement each other well. The difference in the TaL1 results across the two scenarios shows that speaker independent silent speech recognition is a challenging problem. Additionally, we observe a large range in the speaker-wise WERs in the TaL80 test set. Using the Tongue+Lips system, the best speaker

achieves a WER of 32.3% and the worst speaker a WER of 98.3%. Further work can focus on identifying speakers with poor image quality and developing methods to improve their results. There are a number of possibilities that could improve upon these results. For example, we made no attempt to use the audio stream to initialise the silent models. Instead of a flat start, we could bootstrap monophone HMMs with audio alignment. Earlier work showed that substantial improvements can be observed with more complex feature extraction [17, 48]. The results on the TaL80 corpus are speaker independent, but they do make some use of data from target speakers when estimating fMLLR transforms. We observe that using fMLLR can lead to good improvements. Other forms of speaker adaptation are likely to be useful. The multi-speaker systems and fine-tuning results described in the next section further support this idea.

3.2. Speech synthesis

This set of experiments presents results for an articulatory-toacoustic conversion model. As in Section 3.1, we show results in terms of the three modalities available in the TaL corpus. Instead of a speaker-independent scenario, we focus here on **speakerdependent** and **multi-speaker** models for speech recovery from silent modalities. As before, the speaker-dependent model uses days 2-6 from TaL1. In the multi-speaker condition, we use data from 75 speakers in TaL80 and we hold out the rest for additional experiments. For training data, we use all the *aud* utterances available for each speaker, of which 10 utterances are randomly held-out for validation. We use the 24 *xaud* utterances from each speaker for testing.

Feature extraction. The video stream is resampled to match the frame rate of the ultrasound stream using *ffmpeg.* Each video frame is then resized to 72×136 , randomly flipped horizontally, then cropped to 64×128 pixels. Each ultrasound frame is resized to 64×128 . For acoustic features, we use the STRAIGHT vocoder [49] to extract 41-dimensional Mel-cepstral coefficients (MCCs) and

 Table 5. Model structure for the articulatory-to-acoustic model.

Encoder	$\begin{array}{l} \text{3DCNN, } k\ 5\times5\times5, s\ (1,2,2)\ , c\ 32\rightarrow\\ \text{BN-ReLU-Dropout } 0.2\rightarrow\\ \text{Maxpooling, } k\ 1\times2\times2, s\ (1,2,2)\rightarrow\\ \text{3DCNN, } k\ 5\times5\times5, s\ (1,2,2), c\ 64\rightarrow\\ \text{BN-ReLU-Dropout } 0.2\rightarrow\\ \text{Maxpooling, } k\ 1\times2\times2, s\ (1,2,2)\rightarrow\\ \text{3DCNN, } k\ 5\times3\times3, s\ (1,2,2), c\ 128\rightarrow\\ \text{BN-ReLU-Dropout } 0.2\rightarrow\\ \text{Flatten}\rightarrow\text{Fc-512 units}\rightarrow\text{BN-ReLU-Dropout } 0.2\end{array}$
Decoder	2-layer BLSTM, 256 units in each direction \rightarrow Fc-43 units

k and s represent kernel size and stride respectively. c represents the number of output channels. The dimension order is T (time), H (height), W (width). BN represents batch normalization. Fc represents fully connected layer.

1-dimensional fundamental frequency (F0). Fundamental frequency is interpolated at the unvoiced frames and converted to log-F0 and a binary voiced/unvoiced flag. All features are concatenated to form a 43-dimensional acoustic vector. Acoustic features are normalized to have zero mean and unit variance for each speaker. In the multispeaker model, we additionally use a speaker representation in the form of x-vectors [50], extracted with the Kaldi toolkit. For waveform reconstruction, we use the STRAIGHT vocoder.

Model training. We use an encoder-decoder architecture (Table 5). The encoder transforms the tongue or lip frames into 512dimensional vectors. The decoder then predicts acoustic features conditioned on the encoder outputs and speaker x-vectors. When both tongue and lip videos are used as inputs, we adopt separate encoders. The representations from the two encoders are concatenated and sent to the decoder. For multi-speaker training, the model is first pre-trained on all 75 training speakers and then fine-tuned separately on each speaker. We used the Adam optimizer with a learning rates of 10^{-3} and 10^{-4} during pre-training and fine-tuning, respectively.

Results and future work. Results are presented in Table 6. To evaluate the performance of the model, we report results in terms of mel-cepstral distortion (MCD). As a proxy measure of intelligibility, we decode the synthesized samples with an open-source ASR model based on ESPnet [51] and we report results in terms of WER. We opt to use this model rather than the audio model described in section 3.1 as it is not trained on audio data from the target speakers. Decoding the corresponding natural speech with this model achieves a WER of 0.5% and 3.5% for TaL1 and TaL80, respectively. We observe from Table 6 that the performance of the model using different modalities is similar to the results reported in Section 3.1. Using ultrasound tongue images leads to better results than video images of the lips. The best models use both the ultrasound and video data. As with speech recognition, the average performance of the multi-speaker system is worse than that of a speaker-dependent system. We further provide results for fine-tuning the multi-speaker model on the TaL1 data. This is denoted by TaL80+TaL1 in Table 6. Future work will evaluate the intelligibility of synthetic speech with a perceptual test. To this end, a neural vocoder can be used to generate high quality waveforms. More powerful encoders, such as ResNet [52], can be used to process the video inputs. As observed in section 3.1, results vary substantially across speakers. Focusing on identifying speakers whose converted speech has low intelligibility and optimizing for their performance could be an interesting future research direction.

Table 6. Mel-Cepstral Distortion (MCD) and Word Error Rate (WER) for the articulatory-to-acoustic conversion model. In the dataset column, TaL1 denotes results for a speaker-dependent model, TaL80 for a multi-speaker model, and TaL1+TaL80 indicates the multi-speaker model fine-tuned on all available TaL1 data. For TaL80, we show mean \pm standard deviation across all speakers.

Dataset	Data streams	MCD (dB)	WER (%)	
Tal 90	Lips	4.09 ± 0.19	96.0±3.8	
Tallou	Tongue+Lips 3.31 ± 0.15		53.5 ± 11.0 53.5 ± 11.7	
	Lips	3.43	66.3	
TaL1	Tongue Tongue+Lips	2.99 2.84	27.9 17.2	
TaL80+TaL1	Tongue+Lips	2.72	14.0	

Table 7. Accuracy of UltraSync, when pre-trained on out-of-domain child speech therapy data, and when trained on in-domain TaL data.

Prompt type	Tag	TaL1		TaL80	
	Iug	n	accuracy	n	accuracy
	Ou	t-of-doi	nain model		
Read	aud	384	70.6%	2595	71.6%
Read (shared)	xaud	48	79.2%	384	72.9%
Spontaneous	spo	2	100%	16	87.5%
Calibration	cal	18	77.8%	134	80.6%
All		452	71.9%	3129	72.3%
	I	n-doma	in model		
Read	aud	384	98.4%	2595	97.4%
Read (shared)	xaud	48	97.9%	384	98.4%
Spontaneous	spo	2	100%	16	93.8%
Calibration	cal	18	100%	134	98.5%
All		452	98.5%	3129	97.6%

3.3. Automatic Synchronisation

These experiments present results for the automatic synchronisation of ultrasound and audio. The TaL corpus was synchronised at recording time using a hardware synchronisation mechanism which gives the offset between the two signals in milliseconds. We investigate how well we can predict these offsets by exploiting correlations between the two modalities using the UltraSync architecture [14].

We conduct two experiments using in and out-of-domain data. For the first experiment, we use the UltraSync model from [14] which has been trained on **out-of-domain** child speech therapy data. In the second experiment, we train the UltraSync model on **in-domain** TaL data. We reserve days 2, 3, and 4 from TaL1 for training, day 5 for validation, and day 1 and 6 for testing. We also reserve speakers 1-49 from TaL80 for training, 50-65 for validation, and 66-81 for testing. We pool the training data and train a single in-domain model from scratch, following the same architecture and training procedures as [14]. We report all of our results on the same test subset for comparability.

Data preprocessing. We preprocess the data to match the Ultra-



Fig. 3. Pearson's product-moment correlation of speaker-wise results across system pairs. Each cell in this figure indicates the correlation of speaker performance for two systems. WER denotes results from speech recognition systems and MDC denotes results for articulatory-to-acoustic mapping systems.

Sync input size by resampling the audio to 22.05 KHz using *scipy interpolate*, resampling the ultrasound to 24fps using *skimage transform*, and resizing the ultrasound frames to 63×138 pixels.

Experimental setup. UltraSync requires us to specify the range of allowable offsets. In a practical scenario, domain knowledge is utilised to select a suitable candidate list [14]. For these experiments, we consider the observed minimum and maximum offsets given by the hardware sync in the TaL corpus and we set the step size to 45ms. We then calculate the range of allowable offsets as: $[min - (step \times 10)]$, $max + (step \times 10)]$, which renders 25 candidates in total, roughly equal to the 24 candidates in [14].

Results and future work. For each utterance we calculate the discrepancy: disc = prediction - truth, where truth is the hardware offset. A predicted offset is correct if the discrepancy falls within the detectability threshold: -125 < disc < +45 [14]. We report results on utterances with audible content and exclude silent and whispered speech in Table 7. The overall out-of-domain model accuracy is 72.2% while the in-domain accuracy is 97.7%, compared to 82.9% reported by [14]. We attribute the increase in performance to the higher quality of our data, the absence of multiple speakers in our recordings, and the presence of linguistic variety in our utterances compared to child speech therapy data [14]. Future work will explore synchronising the two visual modalities or all three modalities.

4. DISCUSSION

Experimental results show that results can vary depending on tasks or articulatory modalities. In this section, we perform an analysis of speaker performance. We prepare a common evaluation set of all *xaud* utterances available in TaL80 (24 per speaker), which were not seen during training. We collect results separately for each speaker using the corresponding speech recognition and articulatory-to-acoustic mapping systems. We use WER from the TaL80 speaker-independent system in Section 3.1 and MCD from the TaL80 multi-speaker system of Section 3.2. Given these speaker-wise WER/MCD scores, we compute Pearson's productmoment correlation across all system pairs (Figure 3). We observe



Fig. 4. Speaker mean rank for speech recognition and articulatoryto-acoustic mapping using Tongue-only or Lip-only data for TaL80. Speaker identifiers denote speaker number, gender, and accent.

that speaker performance is generally more correlated for intra-task systems than inter-task systems, irrespective of modality used. For example, considering articulatory-to-acoustic mapping (MCD in Figure 3), speaker performance correlates well when using different modalities. That is, speaker-wise MCD results for the systems using tongue-only and lips-only have medium correlation (r = 0.51). But when comparing speaker performance across tasks using the same modality, correlation tends to be lower, either with tongue inputs (r = 0.34) or lip inputs (r = 0.21). We also observe a very high correlation for the articulatory-to-acoustic mapping system using tongue and tongue+lip data (r = 0.98). This suggests that the system may not be making the most use of the video data. This could explain the limited improvements of the tongue+lip system over the tongue system in Section 3.2 when compared with the corresponding system in Section 3.1.

Additionally, we compare speaker performance on systems using only one of the available modalities. To find a comparable score across both tasks, each speaker score (WER or MCD) is replaced by their rank order and then averaged across the two tasks. Figure 4 shows speaker mean ranks for systems using only tongue images or lip images. Although there is low inter-task correlation, Figure 4 identifies speakers that underperfom on both modalities and tasks (e.g. 09fe, 16fe, 27fs, 12me, and 60ms). Similarly, some speakers perform well across the two data streams and tasks (e.g. 70ms, 48ms, and 22me). On the other hand, speaker 63me, for example, achieves good score when using ultrasound data, but scores poorly when using only lip data. This type of analysis can be used in the future to identify speakers with good-quality imaging data.

5. CONCLUSION

We presented the Tongue and Lips corpus (TaL) and a set of benchmark experiments for speech recognition, articulatory-toacoustic mapping, and automatic synchronisation of ultrasound and audio. These experiments illustrate the challenge of developing robust systems that process articulatory imaging data from multiple speakers. TaL is released under the Creative Commons Attribution-NonCommercial 4.0 Generic (CC BY-NC 4.0) licence and is distributed via the UltraSuite Repository: https: //www.ultrax-speech.org/ultrasuite.

6. REFERENCES

- Tanja Schultz, Michael Wand, Thomas Hueber, Dean J Krusienski, Christian Herff, and Jonathan S Brumberg, "Biosignal-based spoken communication: A survey," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2257–2271, 2017.
- [2] Andrew D Scott, Marzena Wylezinska, Malcolm J Birch, and Marc E Miquel, "Speech MRI: morphology and function," *Physica Medica*, vol. 30, no. 6, pp. 604–618, 2014.
- [3] Maureen Stone, "A guide to analysing tongue motion from ultrasound images," *Clinical linguistics & phonetics*, vol. 19, no. 6-7, pp. 455–501, 2005.
- [4] Eleanor Lawson, Jane Stuart-Smith, James M Scobbie, Satsuki Nakai, David Beavan, Fiona Edmonds, Iain Edmonds, Alice Turk, Claire Timmins, Janet M Beck, et al., "Seeing speech: an articulatory web resource for the study of phonetics [website]," 2015, https://www.seeingspeech.ac.uk.
- [5] Manuel Sam Ribeiro, Aciel Eshky, Korin Richmond, and Steve Renals, "Speaker-independent classification of phonetic segments from raw ultrasound in child speech," in *Proc. ICASSP*. IEEE, 2019, pp. 1328–1332.
- [6] Joanne Cleland, James M Scobbie, and Alan A Wrench, "Using ultrasound visual biofeedback to treat persistent primary speech sound disorders," *Clinical linguistics & phonetics*, vol. 29, no. 8-10, pp. 575–597, 2015.
- [7] Joanne Cleland, James M Scobbie, Zoe Roxburgh, Cornelia Heyde, and Alan Wrench, "Enabling new articulatory gestures in children with persistent speech sound disorders using ultrasound visual biofeedback," *Journal of Speech, Language and Hearing Research*, 2018.
- [8] Eleanor Sugden, Susan Lloyd, Jenny Lam, and Joanne Cleland, "Systematic review of ultrasound visual biofeedback in intervention for speech sound disorders," *International journal of language & communication disorders*, vol. 54, no. 5, pp. 705–728, 2019.
- [9] Ian Wilson, Bryan Gick, MG O'Brien, C Shea, and J Archibald, "Ultrasound technology and second language acquisition research," in *Proc. of the 8th Generative Approaches* to Second Language Acquisition Conference (GASLA), 2006, pp. 148–152.
- [10] Bryan Gick, Barbara Bernhardt, Penelope Bacsfalvi, and Ian Wilson, "Ultrasound imaging applications in second language acquisition," *Phonology and second language acquisition*, vol. 36, pp. 315–328, 2008.
- [11] Bruce Denby, Tanja Schultz, Kiyoshi Honda, Thomas Hueber, Jim M Gilbert, and Jonathan S Brumberg, "Silent speech interfaces," *Speech Communication*, vol. 52, no. 4, pp. 270–287, 2010.
- [12] Diandra Fabre, Thomas Hueber, Florent Bocquelet, and Pierre Badin, "Tongue tracking in ultrasound images using eigentongue decomposition and artificial neural networks," in *Proc. Interspeech*, 2015.
- [13] Diandra Fabre, Thomas Hueber, Laurent Girin, Xavier Alameda-Pineda, and Pierre Badin, "Automatic animation of an articulatory tongue model from ultrasound images of the vocal tract," *Speech Communication*, vol. 93, pp. 63–75, 2017.

- [14] Aciel Eshky, Manuel Sam Ribeiro, Korin Richmond, and Steve Renals, "Synchronising audio and ultrasound by learning cross-modal embeddings," *Proc. Interspeech*, 2019.
- [15] Manuel Sam Ribeiro, Aciel Eshky, Korin Richmond, and Steve Renals, "Ultrasound tongue imaging for diarization and alignment of child speech therapy sessions," *Proc. Interspeech*, 2019.
- [16] Thomas Hueber, Elie-Laurent Benaroya, Gérard Chollet, Bruce Denby, Gérard Dreyfus, and Maureen Stone, "Development of a silent speech interface driven by ultrasound and optical images of the tongue and lips," *Speech Communication*, vol. 52, no. 4, pp. 288–300, 2010.
- [17] Yan Ji, Licheng Liu, Hongcui Wang, Zhilei Liu, Zhibin Niu, and Bruce Denby, "Updating the silent speech challenge benchmark with deep learning," *Speech Communication*, vol. 98, pp. 42–50, 2018.
- [18] Bruce Denby and Maureen Stone, "Speech synthesis from real time ultrasound images of the tongue," in *Proc. ICASSP.* IEEE, 2004.
- [19] Tamás Gábor Csapó, Tamás Grósz, Gábor Gosztolya, László Tóth, and Alexandra Markó, "DNN-based ultrasound-tospeech conversion for a silent speech interface," *Proc. Interspeech*, pp. 3672–3676, 2017.
- [20] Thomas Hueber, Elie-Laurent Benaroya, Bruce Denby, and Gérard Chollet, "Statistical mapping between articulatory and acoustic data for an ultrasound-based silent speech interface," in *Proc. Interspeech*, 2011.
- [21] Dagoberto Porras, Alexander Sepúlveda-Sepúlveda, and Tamás Gábor Csapó, "DNN-based acoustic-to-articulatory inversion using ultrasound tongue imaging," in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–8.
- [22] Harry McGurk and John MacDonald, "Hearing lips and seeing voices," *Nature*, vol. 264, no. 5588, pp. 746–748, 1976.
- [23] Joon Son Chung and Andrew Zisserman, "Out of time: automated lip sync in the wild," in Asian conference on computer vision. Springer, 2016, pp. 251–263.
- [24] Adriana Fernandez-Lopez and Federico M Sukno, "Survey on automatic lip-reading in the era of deep learning," *Image and Vision Computing*, vol. 78, pp. 53–72, 2018.
- [25] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, 2018.
- [26] Ariel Ephrat and Shmuel Peleg, "Vid2speech: speech reconstruction from silent video," in *Proc. ICASSP*. IEEE, 2017, pp. 5095–5099.
- [27] Jun Cai, Bruce Denby, Pierre Roussel, Gérard Dreyfus, and Lise Crevier-Buchman, "Recognition and real time performances of a lightweight ultrasound based silent speech interface employing a language model," in *Proc. Interspeech*, 2011.
- [28] Aciel Eshky, Manuel Sam Ribeiro, Joanne Cleland, Korin Richmond, Zoe Roxburgh, James M Scobbie, and Alan A Wrench, "UltraSuite: a repository of ultrasound and acoustic data from child speech therapy sessions," in *Proc. Interspeech*, September 2018.

- [29] Gerasimos Potamianos, Etienne Marcheret, Youssef Mroueh, Vaibhava Goel, Alexandros Koumbaroulis, Argyrios Vartholomaios, and Spyridon Thermos, "Audio and visual modality combination in speech processing applications," in *The Handbook of Multimodal-Multisensor Interfaces: Foundations, User Modeling, and Common Modality Combinations-Volume 1*, pp. 489–543. 2017.
- [30] Iain Matthews, Timothy F Cootes, J Andrew Bangham, Stephen Cox, and Richard Harvey, "Extraction of visual features for lipreading," *IEEE Transactions on Pattern Analysis* and Machine Intelligence, vol. 24, no. 2, pp. 198–213, 2002.
- [31] Kieron Messer, Jiri Matas, Josef Kittler, Juergen Luettin, and Gilbert Maitre, "XM2VTSDB: The extended M2VTS database," in Second international conference on audio and video-based biometric person authentication, 1999, vol. 964, pp. 965–966.
- [32] Martin Cooke, Jon Barker, Stuart Cunningham, and Xu Shao, "An audio-visual corpus for speech perception and automatic speech recognition," *The Journal of the Acoustical Society of America*, vol. 120, no. 5, pp. 2421–2424, 2006.
- [33] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman, "Lip reading sentences in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3444–3453.
- [34] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman, "LRS3-TED: a large-scale dataset for visual speech recognition," arXiv preprint arXiv:1809.00496, 2018.
- [35] Lorenzo Spreafico, Michael Pucher, and Anna Matosova, "Ultrafit: A speaker-friendly headset for ultrasound recordings in speech science," in *Proc. Interspeech*, 2018.
- [36] Articulate Instruments Ltd., Articulate Assistant User Guide: Version 2.11, Articulate Instruments Ltd., Edinburgh, United Kingdom, 2010.
- [37] Articulate Instruments Ltd., *SyncBrightUp Users Manual: Revision 1.10*, Articulate Instruments Ltd., Edinburgh, United Kingdom, 2010.
- [38] Steven Weinberger, "Speech accent archive," *George Mason University*, 2015.
- [39] Grant Fairbanks, *Voice and articulation drillbook*, Harper & Brothers, 1940.
- [40] Institute of Electrical and Electronics Engineers, "IEEE recommended practice for speech quality measurements," *IEEE*, vol. 297, 1969.
- [41] John S Garofolo, "TIMIT acoustic phonetic continuous speech corpus," *Linguistic Data Consortium*, 1993.

- [42] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald, "CSTR VCTK Corpus: English multi-speaker corpus for CSTR Voice Cloning Toolkit," 2019.
- [43] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *Proc. ICASSP.* IEEE, 2015, pp. 5206– 5210.
- [44] ITU-T Recommendation G. 191, "Software tools for speech and audio coding standardization," 2010.
- [45] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., "The Kaldi speech recognition toolkit," in *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2011.
- [46] Shakti P Rath, Daniel Povey, Karel Veselỳ, and Jan Cernockỳ, "Improved feature processing for deep neural networks.," in *Proc. Interspeech*, 2013.
- [47] Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proc. Interspeech*, 2015.
- [48] Eric Tatulli and Thomas Hueber, "Feature extraction using multimodal convolutional neural networks for visual speech recognition," in *Proc. ICASSP*, 2017, pp. 2971–2975.
- [49] Hideki Kawahara, Ikuyo Masuda-Katsuse, and Alain De Cheveigné, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency based F0 extraction: possible role of a repetitive structure in sounds," *Speech Communication*, vol. 27, no. 3–4, pp. 187–207, 1999.
- [50] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. ICASSP.* IEEE, 2018, pp. 5329–5333.
- [51] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique, Yalta Soplin, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [52] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.