



Adaptation of Tacotron2-based Text-To-Speech for Articulatory-to-Acoustic Mapping using Ultrasound Tongue Imaging

Csaba Zainkó¹, László Tóth², Amin Honarmandi Shandiz², Gábor Gosztolya³, Alexandra Markó^{4,5},
Géza Németh¹, Tamás Gábor Csapó^{1,5}

¹Department of Telecommunications and Media Informatics,
Budapest University of Technology and Economics, Budapest, Hungary

²Institute of Informatics, University of Szeged, Hungary

³MTA-SZTE Research Group on Artificial Intelligence, Szeged, Hungary

⁴Department of Applied Linguistics and Phonetics, Eötvös Loránd University, Budapest, Hungary

⁵MTA-ELTE Lendület Lingual Articulation Research Group, Budapest, Hungary

{zainko, nemeth, csapot}@tmit.bme.hu, {tothl, shandiz, ggabor}@inf.u-szeged.hu,
marko.alexandra@btk.elte.hu

Abstract

For articulatory-to-acoustic mapping, typically only limited parallel training data is available, making it impossible to apply fully end-to-end solutions like Tacotron2. In this paper, we experimented with transfer learning and adaptation of a Tacotron2 text-to-speech model to improve the final synthesis quality of ultrasound-based articulatory-to-acoustic mapping with a limited database. We use a multi-speaker pre-trained Tacotron2 TTS model and a pre-trained WaveGlow neural vocoder. The articulatory-to-acoustic conversion contains three steps: 1) from a sequence of ultrasound tongue image recordings, a 3D convolutional neural network predicts the inputs of the pre-trained Tacotron2 model, 2) the Tacotron2 model converts this intermediate representation to an 80-dimensional mel-spectrogram, and 3) the WaveGlow model is applied for final inference. This generated speech contains the timing of the original articulatory data from the ultrasound recording, but the F0 contour and the spectral information is predicted by the Tacotron2 model. The F0 values are independent of the original ultrasound images, but represent the target speaker, as they are inferred from the pre-trained Tacotron2 model. In our experiments, we demonstrated that the synthesized speech quality is more natural with the proposed solutions than with our earlier model.

Index Terms: articulation-to-speech, ultrasound, DNN-TTS

1. Introduction

Articulatory-to-acoustic mapping (AAM) methods aim to synthesize the speech signal directly from articulatory input, as opposed to text-to-speech, when speech is synthesized from the textual input. AAM applies the theory that articulatory movements are directly linked with the acoustic speech signal in the speech production process. A recent potential application of this mapping is a “Silent Speech Interface” (SSI [1, 2, 3]), which has the main idea of recording the soundless articulatory movement, and automatically generating speech from the movement information, while the subject does not produce any sound. Such an SSI system can be highly useful for the speaking impaired (e.g. after laryngectomy or elderly people), and for scenarios where regular speech is not feasible, but the information should be transmitted from the speaker (e.g. extremely noisy environments or military applications).

For the articulatory-to-acoustic mapping, the typical input

can be electromagnetic articulography (EMA) [4, 5], ultrasound tongue imaging (UTI) [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19], permanent magnetic articulography (PMA) [20, 21], surface electromyography (sEMG) [22, 23], Non-Audible Murmur (NAM) [24], electro-optical stomatography [25], impulse radio ultra-wide band (IR-UWB) [26], radar [27] or video of the lip movements [7, 28, 29]. From another aspect, there are two distinct ways of SSI solutions, namely ‘direct synthesis’ and ‘recognition-and-synthesis’ [2]. In the first case, the speech signal is generated without an intermediate step, directly from the articulatory data [4, 5, 6, 8, 9, 11, 12, 14, 15, 16, 20, 22, 23, 24, 28]. In the second case, silent speech recognition (SSR) is applied on the biosignal which extracts the content spoken by the person (i.e. the result of this step is text); this step is then followed by text-to-speech (TTS) synthesis [7, 10, 13, 25, 29, 30]. In the SSR+TTS approach, any information related to speech prosody is lost, whereas it may be kept with direct synthesis. Also, the smaller delay by the direct synthesis approach might enable conversational use.

For the direct conversion, typically, vocoders are used, which synthesize speech from the spectral parameters predicted by the DNNs from the articulatory input. One of the spectral representations that was found to be useful earlier for statistical parametric speech synthesis is Mel-Generalized Cepstrum in Line Spectral Pair form (MGC-LSP) [31, 32]. Since the introduction of WaveNet in 2016 [33], neural vocoders can generate highly natural raw samples of speech, conditioned on mel-spectrogram or other input. One of the most recent types of neural vocoders, WaveGlow [34] is a flow-based network capable of generating high-quality speech from mel-spectrograms. The advantage of the WaveGlow model is that it is relatively simple, yet the synthesis can be done faster than real-time. In [17], we integrated the WaveGlow neural vocoder into ultrasound-based articulatory-to-acoustic conversion.

In the latest years, most TTS solutions apply end-to-end methods, by operating directly on character or phoneme input sequences and producing raw speech signal outputs. One of the most widely used solutions is Tacotron2 [35], which applies a recurrent sequence-to-sequence feature prediction network that maps character embeddings to mel-scale spectrograms, followed by a neural vocoder. The encoder-decoder network, using the attention mechanism, encodes a specific attribute of speech and maps sequences of differing length. In [35], the input char-




Figure 1: Sample ultrasound images from the five sessions.

acters are represented with a learned 512-dimensional embedding, which ensures that traditional text processing is not necessary on the input.

In the field of AAM, according to our knowledge, only a few studies have used fully end-to-end / sequence-to-sequence solutions [36, 37]. Zhang and his colleagues introduced TaLNet, which is based on an encoder-decoder architecture, using the attention mechanism. Both ultrasound and lip are used as the input of AAM, from English speakers of the UltraSuite-TaL database [38]. First, a Tacotron2 model is trained with a large amount of speech data, and after that, transfer learning is applied with the articulatory input. The presented approach was found to be significantly better than earlier baselines. In the study, they also checked the contribution of each articulatory input, and found that the weakest results could be achieved with the lip-only system, followed by ultrasound-only. The combination of ultrasound and lip (TaLNet) was found to be the best, suggesting that these two modalities complement each other well. In another study, by Mira and his colleagues, end-to-end video-to-speech synthesis was proposed, using GANs [37]. The video of the face is translated directly to speech, without an intermediate representation, applying an encoder-decoder architecture. They experimented on various databases and show that the choice of adversarial loss is a key for realistic results.

In this paper, we experiment with transfer learning and adaptation of a Tacotron2 text-to-speech model to improve the final synthesis quality of ultrasound-based articulatory-to-acoustic mapping with a limited database.

2. Methods

2.1. Data

For Tacotron2 and WaveGlow training, we chose 5 male and 6 female Hungarian speakers (altogether 23k sentences, roughly 22 hours) from the PPSD database [39]. This data served as the acoustic-only training material required for the encoder-decoder architecture and the neural vocoder.

For the articulatory data, we used the Hungarian parallel ultrasound and speech dataset that we recorded for earlier studies [16, 17, 40]. We selected a female speaker (speaker048), who was recorded in five sessions (once 209 sentences, and four times 59 sentences). The tongue movement was recorded in midsagittal orientation using the ‘‘Micro’’ ultrasound system of Articulate Instruments Ltd. at 81.67 fps. The speech signal was recorded with a Beyerdynamic TG H56c tan omnidirectional condenser microphone. The ultrasound data and the audio signals were synchronized using the tools provided by Articulate Instruments Ltd. In our experiments, the raw scanline data of the ultrasound was used as input of the networks, after being resized to 64×128 pixels using bicubic interpolation (see samples in Fig. 1), as we found earlier that this reduction does not cause

Conv3D	Conv3D(Transfer_Learning)
Conv3D(30,(5,13,13),Strides=(5,2,2))	Conv3D(30,(5,13,13),Strides=(5,2,2))
Dropout(0.2)	Dropout(0.2)
Conv3D(60,(1,13,13),Strides=(1,2,2))	Conv3D(60,(1,13,13),Strides=(1,2,2))
Dropout(0.2)	Dropout(0.2)
MaxPooling3D(poolsize=(1,2,2))	MaxPooling3D(poolsize=(1,2,2))
Conv3D(70,(1,13,13),Strides=(1,2,2))	Conv3D(70,(1,13,13),Strides=(1,2,2))
Dropout(0.2)	Dropout(0.2)
Conv3D(58,(1,13,13),Strides=(1,1,1))	Conv3D(58,(1,13,13),Strides=(1,1,1))
Dropout(0.2)	Dropout(0.2)
MaxPooling3D(poolsize=(1,2,2))	MaxPooling3D(poolsize=(1,2,2))
Flatten()	Flatten()
Dense(1000)	Dense(1000)
Dropout(0.2)	Dropout(0.2)
Dense(80,activation='linear')	Dense(93, activation='softmax')

Figure 2: The layers of the 3D CNNs in the Keras implementation, along with their most important parameters. Left: baseline 3D CNN for melspectrogram prediction, right: proposed 3D CNN for symbol prediction.

significant information loss [41].

For the Tacotron2 speaker adaptation, speaker048’s data was used (train: 318 sentences, and validation: 40 sentences).

2.2. Ultrasound-to-Melspectrogram using 3D-CNN (baseline)

When we are dealing with image processing as input data, then convolutional neural networks are one of the most popular and effective methods which can extract complex features from data by adding deep layers [42]. In Silent Speech Interface, when we have ultrasound data as input, our input is not only just images but sequences of images which could be considered as a video. Standard CNN considers 2D images to extract features by convolving 2D filters over images. Therefore, to model temporal information, a third dimension has to be considered [43, 44]. Recurrent Neural Networks such as Long Short Term Memory (LSTM) are good examples of combining features extracted from both temporal and spatial parts of data [44]. Using LSTM networks have some drawbacks such as training difficulties, while some variants of these networks were proposed to mitigate this problem, such as quasi-recurrent neural networks [45].

Here we use another variation by adding a third dimension as (2+1)D CNN which shows good performance in video action recognition task [46]. It shows good results when used with ultrasound images and it could be considered as a substitute of CNN+LSTM [18]. In the baseline system of the current study, we apply the same 3D CNN which was used in [18] for predicting 80-dimensional melspectrogram features from ultrasound tongue image input.

This network processed 5 frames of video that were 6 frames apart (6 is the stride parameter of the convolution along the time axis) [18]. Following the concept of (2+1)D convolution, the five frames were first processed only spatially, and then got combined along the time axis just below the uppermost dense layer. Fig. 2 left shows the actual network configuration. The training was performed using the SGD optimizer with 0.06 starting learning rate. It was reduced when a validation MSE has stopped improving by factor 0.5. The batch size was 128. The training objective function was the mean squared error (MSE).

2.3. Ultrasound-to-Symbol using 3D-CNN

In the proposed system, we use the same structure of the 3D CNN as in the baseline system. The difference is in the tar-

get of the network: we predict symbols of Tacotron2 internal representation, having 93 dimensions. At first, we trained with the same methods as the baseline model, but the model was not applicable. We fine-tuned the optimizer, batch size, and other hyperparameters but the model still did not train. Sometimes the accuracy was zero or it learned only the silent symbol and predicted it everywhere. Finally, transfer learning was successful. We reused the baseline 3D-CNN model’s weights at the convolutional layers. All convolutional layers were frozen and only the last two FC layers (with 1000 and 93 neurons) were trained. The weights of these two layers were initialized randomly. Here, cross-entropy is used as the loss function. Because the classes of symbols were not balanced, we used a specific loss function: the loss was weighted with the occurrence of the symbols. We used Adam optimizer and accuracy as a metric. The other parameters of the CNN are the same as the baseline, see Fig. 2 right.

2.3.1. Accuracy and the confusion matrix

The Ultrasound-to-Symbol 3D-CNN model reached 0.68 validation accuracy after 20 epochs (train acc.: 0.83). Early stopping was used with a patience parameter of 7. To improve our Tacotron2 model, the confusion matrix was used to generate augmented training data (see later in Sec. 2.4.3). Fig. 3 shows a simplified version of the confusion matrix (for visualization purposes only – the full matrix involves all 93 symbols: for this figure, we removed the symbols which were not used in the current models and pooled together the short and long versions of the symbols). The values are normalized by rows (target symbols) and converted to percentage values. The first row (on the top) is the most accurate symbol, and the last row (on the bottom) is the least accurate symbol. We expected that the errors are related to articulation, but in Fig. 3 it seems mainly noise-like. The symbols with lower accuracies were some vowels and nasals (e,a,ee,n,m in the figure, /e,ɔ,eɛ,n,m/ in IPA). The symbols with higher accuracies were some less frequent consonants (z,t,y,cs,zs in the figure, /ʒ,tʃ,c,z/ in IPA).

2.4. Symbol-to-melspectrogram using Tacotron2

We used a multi-speaker Tacotron2 model [35] based on the NVIDIA implementation (<https://github.com/NVIDIA/tacotron2>). The speakers’ IDs are coded as a one-hot vector and added to the inputs of the LSTM cells both in the encoder and decoder. The model was trained by all 11 speakers of the PPSD database [39] at the same time. The order of all speakers’ sentences was randomized. The input of the Tacotron2 is a sequence of symbols. Because Hungarian is an almost phonetic language, we used a mixed collection of letters and phonemes. The symbols of the input sequence follow the phonemes of the sentences, but we did not use allophones or other detailed discrimination. Only the long–short property is used to encode durational differences. The phonemes are represented with their approximate letter: the lowercase letters show the short phonemes, the capital letters indicate the long phonemes.

This multi-speaker model was trained during 156k iterations on a single NVIDIA Titan Xp. The sample rate of the sound was 22 050 Hz, the window size was 1024 and the hop length was 256. We used 80 mel channels between 0 Hz and 8000 Hz to keep compatibility with the WaveGlow model. The encoder’s symbols embedding and embedding dimension was also 512. The decoder’s RNN dimensions were 1024.

Our goal was to use our pre-trained Tacotron2 model (orig-

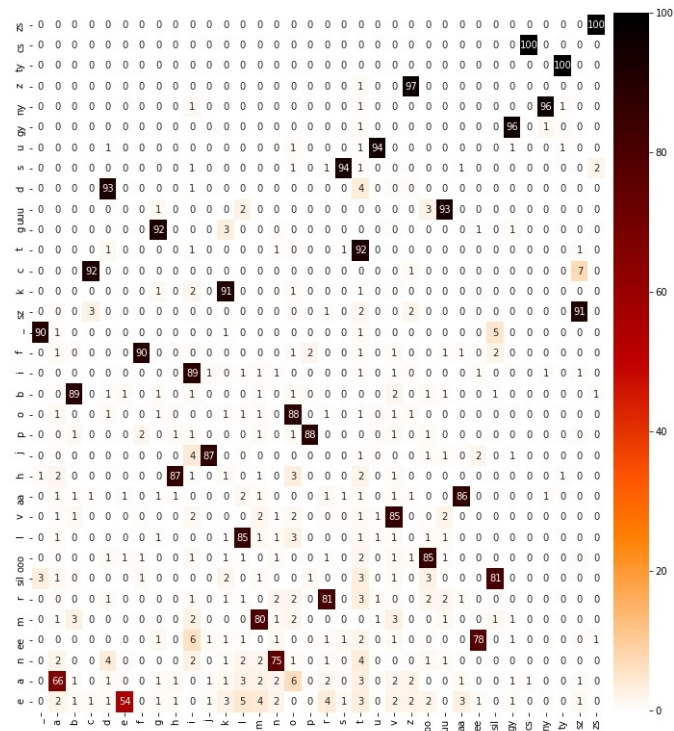


Figure 3: *Simplified confusion matrix of the proposed Ultrasound-to-Symbol 3D-CNN. The values are normalized and showed in percentages. Rows: target, columns: predicted.*

inally developed for TTS) without modification, therefore we made only some fine-tuning for AAM purposes. The ultrasound image sequence does not contain F0-related information, but it contains the timing of speech. Basically, the Tacotron2 does not handle timing information of a sentence, it can generate that via an attention mechanism. Fig. 4 top shows an example for the connection between the steps of the encoder and decoder with this initial Tacotron2 system. This sentence encoder contains 16 symbols plus two padding symbols at the borders of the sentence. The model generated 134 decoder frames. In this model, one frame is about 11.6ms, so this sentence was about 1.6s long. Clearly, the timings are not modeled well here.

2.4.1. Time-synchronous Tacotron2 system

In order to use the proper timing of the input sequence, we generated a new training set from the original 11 speakers’ dataset. The input symbols were repeated accordingly to the real duration of a phone. The repeating number was calculated from the ultrasound frame rate (81.67 fps). For example, at a 98ms long phone, the symbol was repeated 8 times. The attention mechanism adapted to the synchronized input during the fine-tuning. It required 7.5k iterations.

2.4.2. Proposed system #1

The speaker in the ultrasound dataset (speaker048) is independent of the 11 speakers of the training set of Tacotron2. The next step was fine-tuning to the new speaker. We chose a female speaker from the 11 others, and at the tuning, her speakerID one-hot vector was used. At this step, 84 iterations resulted in the smallest validation error. In the first proposed system,

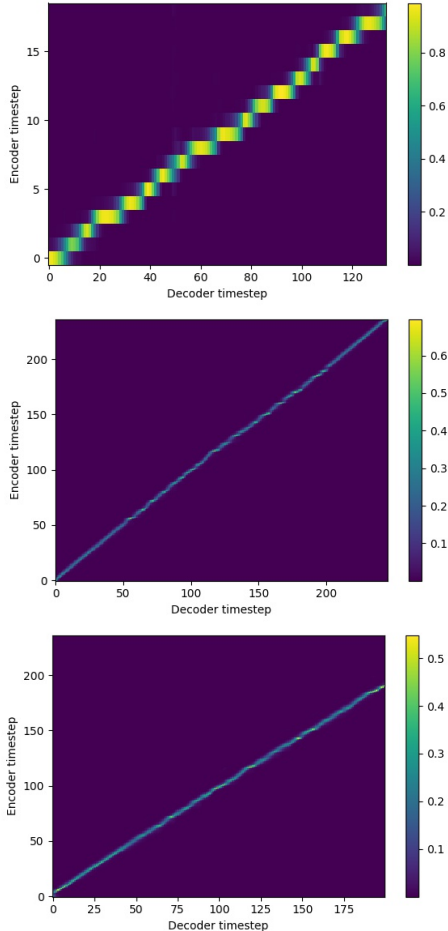


Figure 4: Examples for the connection between the steps of the encoder and decoder. Top: Tacotron2 without timing information. Middle: Tacotron2 with timing information (Proposed #1). Bottom: Tacotron2 with timing information and with data augmentation (Proposed #2).

this model was used. Fig. 4 middle shows the proper timing of the generated speech. The input of that sentence contains 237 symbols, and the system generated 246 output frames. The difference comes from the uncertainty of the end decision of the decoder. The figure also shows the Tacotron2 can tolerate some symbol errors, i.e. the line is not perfectly straight; there are some small steps, where the decoder ignores some input symbols.

2.4.3. Proposed system #2

Our experience was that Tacotron2 can tolerate some mistakes in the prediction of the 3D-CNN model (Sec. 2.3), but these mistakes cause audible distortion during the final synthesis. The distribution of the wrong predictions can be characterized by the confusion matrix (Sec. 2.3.1) of the 3D-CNN network. It is not accurate because it does not contain the position information of the mistakes, but it is suitable to generate similar training data for fine-tuning the Tacotron2 model. With the distribution of the symbol’s error, we modified the 11 speakers training set. The symbol changing was based on the distribution but it was randomized. For every sentence, 20 different versions were generated. The output mel-spectrograms were not changed. 4.3k

iterations provided the lowest validation error. Fig. 4 bottom shows the tuned model’s connection between the encoder and decoder. There are two differences compared to the middle sub-figure. The number of the encoder steps remained the same, but there are fewer decoder steps. The decoder learned to ignore the different types of silence symbols (pad, sil, start_sil, end_sil) which were mixed in the predicted symbol sequence. The other difference is that the line is smoother. It shows that a decoder step connects more encoder steps and the model can combine the information of good and bad symbols.

After that we also repeated the tuning to the speaker from the ultrasound dataset. Here we also generate modified training data with the phoneme errors. The procedure was the same as at the multi-speaker case. At this second step, 182 iterations were required. We used this model in the second proposed system.

2.5. Melspectrogram-to-speech with a neural vocoder

Similarly to the original WaveGlow paper [34], 80 bins were used for mel-spectrogram using librosa mel-filter defaults (i.e. each bin is normalized by the filter length and the scale is the same as in HTK, Hidden Markov Model Toolkit). FFT size and window size were both 1024 samples. For hop size, we use the base 256 samples. This 80-dimensional mel-spectrogram served as the training target of the Tacotron2 network. A WaveGlow model was trained with the Hungarian data (WaveGlow-HU). This latter training was done on a server with eight V100 GPUs, altogether for 635k iterations. In the synthesis phase, an interpolation in time was not necessary, different from [17]. The ultrasound frame rate was 270 samples, but the differences were compensated by the Tacotron2 model, the output frame rate of the model was 256 samples which is the same as the WaveGlow’s hop size. Finally, the synthesized speech is the result of the inference with the trained WaveGlow-HU model conditioned on the mel-spectrogram input [34].

3. Experiments and Results

After training the above models, we synthesized sentences from the test part of the ultrasound dataset. These sentences have not been used during the training process, neither in the Ultrasound-to-Symbol model, nor in the Tacotron2 training and tuning process. The domain of the texts is also independent of the training and validation dataset: it contains the Hungarian version of ‘The North Wind and the Sun’.

3.1. Subjective listening test

In order to determine which proposed version is closer to natural speech, we conducted an online MUSHRA-like test [47]. Our aim was to compare the natural sentences with the synthesized sentences of the baseline, the proposed approaches and a lower anchor system (the latter having constant F0 and 2D CNN predicted MGC-LSP, from [17]). In the test, the listeners had to rate the naturalness of each stimulus in a randomized order relative to the reference (which was the natural sentence), from 0 (very unnatural) to 100 (very natural). We chose nine sentences from the test set of the target speaker. The variants appeared in randomized order (different for each listener). The samples can be found at http://smartlab.tmit.bme.hu/ssw11_tacotron2.

Each sentence was rated by 23 native Hungarian speakers (11 females, 12 males; 14–47 years old), in a silent environment. On average, the test took 10 minutes to complete. Fig. 5 shows the average naturalness scores for the tested approaches.

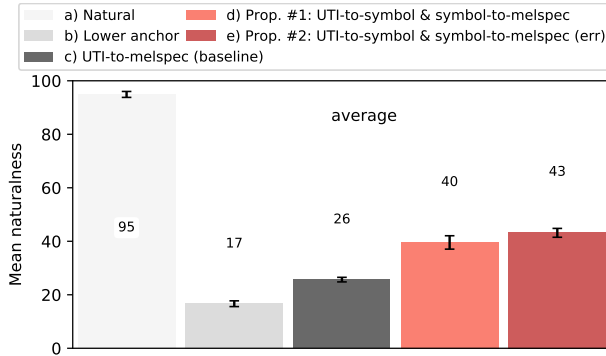


Figure 5: Results of the subjective evaluation with respect to naturalness. The error bars show the 95% confidence intervals.

The lower anchor received the weakest scores, followed by the baseline, and the proposed approaches. To check the statistical significances, we conducted Mann-Whitney-Wilcoxon ranksum tests with a 95% confidence level. Based on this, both proposed variants were evaluated as significantly more natural than the baseline. The listeners noted the difference between the two proposed versions: proposed#1, the one with standard training (Sec. 2.4.2) was rated as 40%, while proposed #2, the one with additional error training (Sec. 2.4.3) was rated as 43% – but this difference is not statistically significant.

As a summary of the listening test, we can conclude that splitting the ultrasound-to-speech prediction task into three parts increased the naturalness, mostly because of the Tacotron2 component which could be trained with a large amount of speech data, and transfer learning / adaptation was possible to the target speaker.

4. Discussion

In Sec. 1, we noted that currently only a few sequence-to-sequence / fully end-to-end solutions are available for articulatory-to-acoustic mapping [36, 37]. Our proposed solution has the following similarities and differences. Mira and his colleagues use the video of the face as input [37], Zhang and his colleagues use both ultrasound and lip video input [36], whereas in our study we use ultrasound tongue image input. As the three studies apply different databases, the results are not directly comparable. In [37], GANs are used with specific adversarial loss, whereas we apply 3D CNN to model the spatial and temporal dependencies of the articulatory and acoustic data. Similarly to [36], we apply Tacotron2 as the encoder-decoder network, but we extend the basic training with additional data augmentation, which includes the wrong predictions from the confusion matrix of the UTI-to-symbol prediction network. By using the symbols as intermediate representation, our solution is closer to the ‘recognition-and-synthesis’ type of SSIs.

5. Conclusions

In this paper, we experimented with transfer learning and adaptation of a Tacotron2 text-to-speech model to improve the final synthesis quality of ultrasound-based articulatory-to-acoustic mapping with a limited database (roughly 200 sentences). We used a Hungarian multi-speaker pre-trained Tacotron2 TTS model and a pre-trained WaveGlow neural vocoder (both trained on 11 speakers’s data, altogether 23k sentences, roughly 22 hours of speech). The proposed articulatory-to-acoustic

conversion framework is a fully end-to-end solution, including an encoder-decoder architecture and attention mechanism, and contains three steps: 1) from a sequence of ultrasound tongue image recordings, a 3D convolution neural network predicts the 93-dimensional embedding inputs of the pre-trained Tacotron2 model, 2) the Tacotron2 model converts this intermediate representation to a 80-dimensional mel-spectrogram, and 3) the WaveGlow model is applied for final inference. We demonstrated that the synthesized speech quality is significantly more natural with the proposed solutions than with our earlier model.

The code is accessible at <https://github.com/BME-SmartLab/UTI-to-STFT-Tacotron2>.

6. Acknowledgements

The research was partly supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 825619 (AI4EU), by the National Research Development and Innovation Office of Hungary (FK 124584 and PD 127915 grants; APH-ALARM / 2019-2.1.2-NEMZ-2020-00012 project) and through the Artificial Intelligence National Laboratory Programme. The Titan X GPU used was donated by NVIDIA Corporation. We would like to thank the subjects for participating in the listening test.

7. References

- [1] B. Denby, T. Schultz, K. Honda, T. Hueber, J. M. Gilbert, and J. S. Brumberg, “Silent speech interfaces,” *Speech Communication*, vol. 52, no. 4, pp. 270–287, 2010.
- [2] T. Schultz, M. Wand, T. Hueber, D. J. Krusienski, C. Herff, and J. S. Brumberg, “Biosignal-Based Spoken Communication: A Survey,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2257–2271, dec 2017.
- [3] J. A. Gonzalez-Lopez, A. Gomez-Alanis, J. M. Martin Donas, J. L. Perez-Cordoba, and A. M. Gomez, “Silent Speech Interfaces for Speech Restoration: A Review,” *IEEE Access*, vol. 8, pp. 177 995–178 021, sep 2020.
- [4] B. Cao, M. Kim, J. R. Wang, J. Van Santen, T. Mau, and J. Wang, “Articulation-to-Speech Synthesis Using Articulatory Flesh Point Sensors’ Orientation Information,” in *Proc. Interspeech*, Hyderabad, India, 2018, pp. 3152–3156.
- [5] F. Taguchi and T. Kaburagi, “Articulatory-to-speech conversion using bi-directional long short-term memory,” in *Proc. Interspeech*, Hyderabad, India, 2018, pp. 2499–2503.
- [6] B. Denby and M. Stone, “Speech synthesis from real time ultrasound images of the tongue,” in *Proc. ICASSP*, Montreal, Quebec, Canada, 2004, pp. 685–688.
- [7] T. Hueber, E.-L. Benaroya, G. Chollet, G. Dreyfus, and M. 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.
- [8] T. Hueber, E.-L. Benaroya, B. Denby, and G. Chollet, “Statistical Mapping Between Articulatory and Acoustic Data for an Ultrasound-Based Silent Speech Interface,” in *Proc. Interspeech*, Florence, Italy, 2011, pp. 593–596.
- [9] A. Jaumard-Hakoun, K. Xu, C. Leboullenger, P. Roussel-Ragot, and B. Denby, “An Articulatory-Based Singing Voice Synthesis Using Tongue and Lips Imaging,” in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 1467–1471.
- [10] E. Tatulli and T. Hueber, “Feature extraction using multimodal convolutional neural networks for visual speech recognition,” in *Proc. ICASSP*, New Orleans, LA, USA, 2017, pp. 2971–2975.
- [11] T. G. Csapó, T. Grósz, G. Gosztolya, L. Tóth, and A. Markó, “DNN-Based Ultrasound-to-Speech Conversion for a Silent Speech Interface,” in *Proc. Interspeech*, Stockholm, Sweden, 2017, pp. 3672–3676.

- [12] T. Grósz, G. Gosztolya, L. Tóth, T. G. Csapó, and A. Markó, “F0 Estimation for DNN-Based Ultrasound Silent Speech Interfaces,” in *Proc. ICASSP*, Calgary, Canada, 2018, pp. 291–295.
- [13] L. Tóth, G. Gosztolya, T. Grósz, A. Markó, and T. G. Csapó, “Multi-Task Learning of Phonetic Labels and Speech Synthesis Parameters for Ultrasound-Based Silent Speech Interfaces,” in *Proc. Interspeech*, Hyderabad, India, 2018, pp. 3172–3176.
- [14] E. Moliner and T. G. Csapó, “Ultrasound-based silent speech interface using convolutional and recurrent neural networks,” *Acta Acustica united with Acustica*, vol. 105, no. 4, pp. 587–590, 2019.
- [15] G. Gosztolya, Á. Pintér, L. Tóth, T. Grósz, A. Markó, and T. G. Csapó, “Autoencoder-Based Articulatory-to-Acoustic Mapping for Ultrasound Silent Speech Interfaces,” in *International Joint Conference on Neural Networks*, 2019.
- [16] T. G. Csapó, M. S. Al-Radhi, G. Németh, G. Gosztolya, T. Grósz, L. Tóth, and A. Markó, “Ultrasound-based Silent Speech Interface Built on a Continuous Vocoder,” in *Proc. Interspeech*, Graz, Austria, 2019, pp. 894–898.
- [17] T. G. Csapó, C. Zainkó, L. Tóth, G. Gosztolya, and A. Markó, “Ultrasound-based Articulatory-to-Acoustic Mapping with WaveGlow Speech Synthesis,” in *Proc. Interspeech*, 2020, pp. 2727–2731.
- [18] L. Tóth and A. H. Shandiz, “3D Convolutional Neural Networks for Ultrasound-Based Silent Speech Interfaces,” in *Proc. ICAISC*, Zakopane, Poland, 2020.
- [19] A. H. Shandiz, L. Tóth, G. Gosztolya, A. Markó, and T. G. Csapó, “Improving Neural Silent Speech Interface Models by Adversarial Training,” in *2nd International Conference on Artificial Intelligence and Computer Vision (AICV2021)*, 2021.
- [20] J. A. Gonzalez, L. A. Cheah, A. M. Gomez, P. D. Green, J. M. Gilbert, S. R. Ell, R. K. Moore, and E. Holdsworth, “Direct Speech Reconstruction From Articulatory Sensor Data by Machine Learning,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2362–2374, dec 2017.
- [21] J. A. Gonzalez-Lopez, M. Gonzalez-Atienza, A. Gomez-Alanis, J. L. Perez-Cordoba, and P. D. Green, “Multi-view Temporal Alignment for Non-parallel Articulatory-to-Acoustic Speech Synthesis,” in *Proc. IberSPEECH*, 2021, pp. 230–234.
- [22] M. Janke and L. Diener, “EMG-to-Speech: Direct Generation of Speech From Facial Electromyographic Signals,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 12, pp. 2375–2385, dec 2017.
- [23] L. Diener, G. Felsch, M. Angrick, and T. Schultz, “Session-Independent Array-Based EMG-to-Speech Conversion using Convolutional Neural Networks,” in *13th ITG Conference on Speech Communication*, 2018.
- [24] N. Shah, N. Shah, and H. Patil, “Effectiveness of Generative Adversarial Network for Non-Audible Murmur-to-Whisper Speech Conversion,” in *Proc. Interspeech*, Hyderabad, India, 2018, pp. 3157–3161.
- [25] S. Stone and P. Birkholz, “Silent-speech command word recognition using electro-optical stomatography,” in *Proc. Interspeech*, San Francisco, CA, USA, 2016, pp. 2350–2351.
- [26] Y. H. Shin and J. Seo, “Towards contactless silent speech recognition based on detection of active and visible articulators using IR-UWB radar,” *Sensors*, vol. 16, no. 11, 2016.
- [27] P. A. Digehsara, C. Wagner, P. Schaffer, M. Bärhold, S. Stone, D. Plettemeier, and P. Birkholz, “On the optimal set of features and robustness of classifiers in radar-based silent phoneme recognition,” in *Proc. ESSV*, online, 2021.
- [28] A. Ephrat and S. Peleg, “Vid2speech: Speech Reconstruction from Silent Video,” in *Proc. ICASSP*, New Orleans, LA, USA, 2017, pp. 5095–5099.
- [29] K. Sun, C. Yu, W. Shi, L. Liu, and Y. Shi, “Lip-Interact: Improving Mobile Device Interaction with Silent Speech Commands,” in *UIST 2018 - Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, Berlin, Germany, 2018, pp. 581–593.
- [30] F. V. Arthur and T. G. Csapó, “Towards a practical lip-to-speech conversion system using deep neural networks and mobile application frontend,” in *2nd International Conference on Artificial Intelligence and Computer Vision (AICV2021)*, 2021.
- [31] T. G. Csapó, G. Németh, and M. Cernak, “Residual-Based Excitation with Continuous F0 Modeling in HMM-Based Speech Synthesis,” in *Lecture Notes in Artificial Intelligence*, A.-H. Dediu, C. Martín-Vide, and K. Vicsi, Eds. Budapest, Hungary: Springer International Publishing, 2015, vol. 9449, pp. 27–38.
- [32] T. G. Csapó, G. Németh, M. Cernak, and P. N. Garner, “Modeling Unvoiced Sounds In Statistical Parametric Speech Synthesis with a Continuous Vocoder,” in *Proc. EUSIPCO*, Budapest, Hungary, 2016, pp. 1338–1342.
- [33] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “WaveNet: A Generative Model for Raw Audio,” *CoRR*, vol. abs/1609.0, 2016.
- [34] R. Prenger, R. Valle, and B. Catanzaro, “Waveglow: A Flow-based Generative Network for Speech Synthesis,” in *Proc. ICASSP*, Brighton, UK, 2019, pp. 3617–3621.
- [35] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, “Natural TTS Synthesis by Conditioning Wavenet on MEL Spectrogram Predictions,” in *Proc. ICASSP*, Calgary, Canada, 2018, pp. 4779–4783.
- [36] J.-X. Zhang, K. Richmond, Zhen-Hua-Ling, and L.-R. Dai, “TaLNet: Voice Reconstruction from Tongue and Lip Articulation with Transfer Learning from Text-to-Speech Synthesis,” in *Proc. AAAI*, 2021.
- [37] R. Mira, K. Vougioukas, P. Ma, S. Petridis, B. W. Schuller, and M. Pantic, “End-to-End Video-To-Speech Synthesis using Generative Adversarial Networks,” apr 2021.
- [38] M. S. Ribeiro, J. Sanger, J.-X. X. Zhang, A. Eshky, A. Wrench, K. Richmond, and S. Renals, “TaL: a synchronised multi-speaker corpus of ultrasound tongue imaging, audio, and lip videos,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*, Shenzhen, China, 2021, pp. 1109–1116.
- [39] G. Olasz, “Precíziós, párhuzamos magyar beszédadatbázis fejlesztése és szolgáltatásai [Development and services of a Hungarian precisely labeled and segmented, parallel speech database] (in Hungarian),” *Beszédkutatás 2013 [Speech Research 2013]*, pp. 261–270, 2013.
- [40] G. Gosztolya, T. Grósz, L. Tóth, A. Markó, and T. G. Csapó, “Applying DNN Adaptation to Reduce the Session Dependency of Ultrasound Tongue Imaging-Based Silent Speech Interfaces,” *Acta Polytechnica Hungarica*, vol. 17, no. 7, pp. 109–124, 2020.
- [41] T. G. Csapó, G. Gosztolya, L. Tóth, A. H. Shandiz, and A. Markó, “Optimizing the Ultrasound Tongue Image Representation for Residual Network-based Articulatory-to-Acoustic Mapping,” *submitted to Multimedia Tools and Applications*, 2021.
- [42] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [43] S. Ji, W. Xu, M. Yang, and K. Yu, “3D convolutional neural networks for human action recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, 2012.
- [44] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, nov 1997.
- [45] J. Bradbury, S. Merity, C. Xiong, and R. Socher, “Quasi-recurrent neural networks,” *arXiv preprint arXiv:1611.01576*, 2016.
- [46] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, “A closer look at spatiotemporal convolutions for action recognition,” in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 6450–6459.
- [47] “ITU-R Recommendation BS.1534: Method for the subjective assessment of intermediate audio quality,” 2001.