



Adaptation of Tongue Ultrasound-Based Silent Speech Interfaces Using Spatial Transformer Networks

László Tóth¹, Amin Honarmandi Shandiz¹, Gábor Gosztolya^{1,2}, Tamás Gábor Csapó³

¹Institute of Informatics, University of Szeged, Hungary

²ELRN-SZTE Research Group on Artificial Intelligence, Szeged, Hungary

³Department of Telecommunications and Media Informatics,

Budapest University of Technology and Economics, Budapest, Hungary

{ tothl, shandiz, ggabor } @ inf.u-szeged.hu, csapot@tmit.bme.hu

Abstract

Thanks to the latest deep learning algorithms, silent speech interfaces (SSI) are now able to synthesize intelligible speech from articulatory movement data under certain conditions. However, the resulting models are rather speaker-specific, making a quick switch between users troublesome. Even for the same speaker, these models perform poorly cross-session, i.e. after dismounting and re-mounting the recording equipment. To aid quick speaker and session adaptation of ultrasound tongue imaging-based SSI models, we extend our deep networks with a spatial transformer network (STN) module, capable of performing an affine transformation on the input images. Although the STN part takes up only about 10% of the network, our experiments show that adapting just the STN module might allow to reduce MSE by 88% on the average, compared to retraining the whole network. The improvement is even larger (around 92%) when adapting the network to different recording sessions from the same speaker.

Index Terms: silent speech interface, ultrasound tongue imaging, speaker adaptation, spatial transformer network

1. Introduction

The goal of articulatory-to-acoustic mapping (AAM) is to synthesize speech from data capturing the movement of the articulatory organs. This is the purpose of the giving the background to the development of 'Silent Speech Interface' systems (SSI [1, 2]). Ideally, these interfaces would record the articulation and synthesize speech based on the movement of the organs – without the user of the device actually producing any sound. The typical input of AAM can be a video of the lip movements [3, 4, 5, 6, 7, 8], ultrasound tongue imaging (UTI) [3, 9, 10, 11, 12, 13, 14, 15, 16, 17], or several other modalities (e.g., MRI, EMA, PMA, EOS, radar, multimodal, etc.). All of the articulatory tracking devices are highly sensitive to 1) the alignment of the recording equipment across sessions, 2) the actual speaker's anatomy. For example, in the case of ultrasound recordings, the probe fixing headset has to be re-mounted onto the speaker for each recording session. This inevitably causes the recorded ultrasound videos to become misaligned between each recording session [18]. Moreover, there are large individual differences across speakers, so even a system trained on the data of several speakers may still perform poorly for a new speaker.

There have already been several cross-session and cross-speaker studies, of which we mention only those related to imaging. To handle the session dependency of UTI-based synthesis, Gosztolya et al. used data from different sessions [19]. Ribeiro et al. reported that, for a speaker-independent system, unsupervised model adaptation can improve the results

for silent speech [20]. In a multi-speaker framework, Shandiz et al. experimented with x-vectors features extracted from speakers, leading to a marginal improvement in the spectral estimation step [17]. Zhang et al. evaluated UTI and lip video based unconstrained multi-speaker voice recovery with a transfer learning strategy and encoder-decoder architecture [21]. There have been more studies on multi-speaker lip-to-speech synthesis [5, 6, 7, 8]. One of the first papers that could produce intelligible speech for unseen speakers was based on WGAN with new additional critics and losses [5]. Another study proposed speaker disentanglement by inputting speaker identities or embeddings to the DNN [6]. An end-to-end ResNet-based model was claimed to outscore previous approaches on unseen speakers [7], while Wang et al. applied cross-modal knowledge transfer and voice conversion to generate good quality speech with high naturalness [8].

Most of the above approaches hope to solve speaker sensitivity simply by acquiring articulatory training data from a large quantity of speakers. In this study, we experiment with a direct adaptation of an UTI-based SSI network to the actual speaker or session. To avoid the need for a full retraining, we extend our network with a spatial transformer network (STN) module and retrain only this module during the adaptation step. The STN learns an affine transformation on the input images, and our assumption is that this transformation should mostly be able to compensate for the misalignment of the recording device, and to a certain extent also for the inter-speaker differences.

2. The UTI-to-Speech framework

Approaches for articulatory-to-acoustic mapping similar to ours were described several times [15, 17, 22], so we just give a brief overview here (see also Fig. 1). The input to our system is a sequence of ultrasound tongue images, which record the movement of the articulators at a relatively high frame rate. Our goal is to estimate the speech signal produced during articulation, but instead of producing a speech signal, we estimate only a mel-spectrogram at the network output, and we apply a neural vocoder to convert the spectrogram to speech. The main advantage of this approach is that the mel-spectrum is a very dense representation, which is easier to estimate than the speech signal itself, and we can apply large pre-trained networks for the synthesis step, such as WaveGlow [23]. Our training data consists of precisely time-aligned pairs of ultrasound videos and speech signals, so the ultrasound-to-spectrogram mapping can be performed on a simple frame-by-frame basis, converting each ultrasound image to a spectral vector [11, 15]. While this simple arrangement already performs reasonably well, significant improvement can be achieved by involving the input context, that is, by using a block of video frames as input instead of

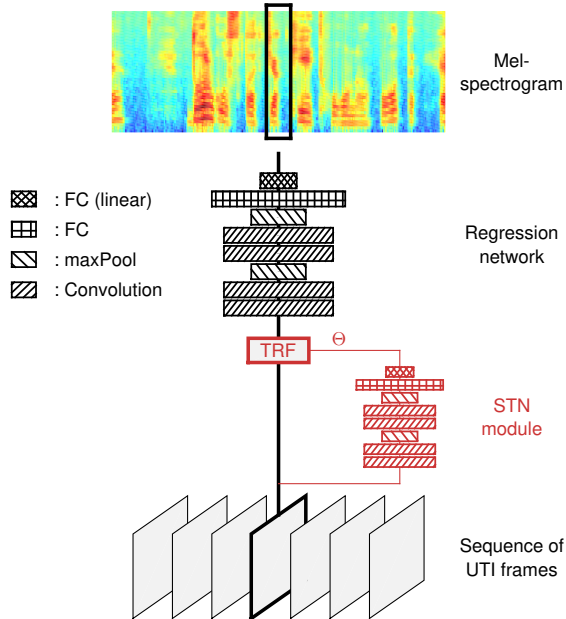


Figure 1: Illustration of the network architecture used, where FC denotes the Fully-Connected layer.

just one image. Several network architectures have been proposed to process 3D blocks of input data, for video processing in general [24, 25, 26], and for ultrasound input in particular [12, 14, 27, 16, 28]. In the experimental section we will experiment both with 2D and 3D Convolutional Neural Networks (CNNs) for the mapping task. The problem could also be addressed even in the lack of aligned training data using encoder-decoder networks [20, 21] or video transformers [29, 30].

3. The Spatial Transformer Network

The concept of the Spatial Transformer Network (STN) was motivated by the fact that in image recognition the actual input is frequently shifted, rotated or scaled compared to the training data, due to changes in camera angle, distance, and other factors [31]. Although CNNs are somewhat invariant to translations, their overall flexibility would greatly improve by introducing a dynamic mechanism that can spatially transform the input image by an appropriate transformation before classification. The STN is a network module that can be inserted into a CNN architecture at any point, but typically it is applied between the input and the first network layer. In this arrangement, the STN performs an affine transformation on the input image, and returns a transformed image of the same size. As it manipulates only the input, it can be combined with practically any type of classification or regression network. The affine transform defined by the STN is quite powerful and includes translation, scaling, rotation, shearing and cropping as special cases.

The STN consists of three main parts, namely the localization network, the grid generator and the sampler [31]. The grid generator and the sampler together are responsible for executing the affine transformation defined by the parameter θ , which consists of 6 components in the case of a 2D input (see Fig. 1).

The grid generator and the sampler are differentiable. This is vital for propagating the error back to the third component, the localization network, which learns to estimate the θ parameters, conditioned on the actual input image. The localization network can take any form, but its uppermost layer must be a regression layer to produce the θ values.




Figure 2: An example UTI image, before (left) and after (right) STN.

While we explained the STN concept assuming 2D images, the whole idea can be naturally extended to 3D data blocks [31, 32]. The 3D variant has already been used in visual speech recognition (e.g., lip reading) by Yu and Wang [33]. We discuss the options for a 3D extension in Section 4.3.

Although the original concept uses the STN to decrease the variance of the training data by transforming the input images to a canonical, expected pose, it may also be useful in handling a domain mismatch between the training and testing conditions. It was applied for domain adaptation in various image processing situations [34, 35], and here we evaluate its efficiency for speaker and session adaptation in UTI-to-speech conversion.

4. Experimental Set-Up

4.1. Data Acquisition and Preprocessing

Several Hungarian male and female subjects with normal speaking abilities were recorded while reading sentences aloud (altogether 209 sentences each), of which two female and two male speakers were selected for the speaker adaptation experiments (with speaker IDs 048, 049, 102 and 103). The average duration of the recordings was about 15 minutes per speaker. For the cross-session experiments one female speaker from the 4 speakers (speaker 048) was asked to record 4 additional sessions on different days (obviously, with dismounting and remounting the UTI recording device between the sessions). These 4 additional sessions are shorter than the main session of the given speaker, with an average duration of 3.5 minutes. All sessions were divided into training, validation and test partitions randomly, in a 85-10-5 ratio. The speech signal and the ultrasound articulatory data were recorded in parallel, and were synchronized using the software provided with the UTI equipment, Figure 3. The tongue movement was recorded in a midsagittal orientation using the “Micro” ultrasound system of Articulate Instruments Ltd. at 81.67 fps. The speech signal was recorded with an omnidirectional condenser microphone. For more details on the recording set-up see [10]. In the experiments we used the raw scanline data of the ultrasound as the training input, after being resized to 64×128 pixels. The intensity range of the images was min-max normalized to the [-1, 1] interval. The speech signals were converted to a mel-spectrogram with 80 spectral components, according to the requirements of WaveGlow. These 80-component spectral vectors served as the training targets for our neural networks, after standardization to zero mean and unit variance.

4.2. Network configuration

In the first set of experiments we applied a simple 2D convolutional (2D-CNN) network that transforms one ultrasound image to one spectral vector, and the localization network of the STN also was a 2D-CNN. Although this approach is suboptimal, it may already show whether the STN is suitable to perform domain adaptation, while visualization and interpretation is easier in 2D. Fig. 1 illustrates this network arrangement, while Fig. 2 demonstrates the effect of STN on an actual ultrasound image. The 2D-CNN network had a quite simple and traditional structure. It consisted of 4 convolutional layers with 30-60-90-120 filters, with a MaxPooling layer after every second layer. The convolutional processing was followed by a fully connected layer of 300 neurons, and the linear output layer of 80 neurons. All hidden layers applied the Swish nonlinearity, and overfitting was minimized by placing dropout layers ($p=0.2$) after each processing layer. We applied the mean squared error (MSE) loss function, which was minimized using the Adam optimizer with a batch size of 100 and a learning rate of $2 \cdot 10^{-3}$. Training was halted using early stopping on the validation set.

The input of the STN module is the same image as that of the spectral regression module, and it also has to perform regression. Hence we used exactly the same 2D-CNN architecture for the STN as for the main network. However, as the STN has to estimate only 6 parameters (the θ vector) instead of 80, we configured it to be much smaller than the regression network. The number of free parameters in the STN module was only about 10% of that of the whole network.

4.3. Extension to 3D

The simple framework presented above can be significantly improved by extending the input from just one image to a sequence of video frames. Such 3D blocks of input data can be processed by various network types, such as by combining 2D-CNN and LSTM layers, by 3D-CNN models, or by using Convolutional LSTM layers [12, 14, 28, 27, 36]. Here, we applied the 3D-CNN architecture introduced by Tóth et al. [27, 17]. This network had the same global architecture as the 2D-CNN presented above. However, the convolutions were extended to the time axis as well, in a somewhat special arrangement. The first Conv3D layer applied a large stride along the time axis, dividing the sequence of 25 input images into 5 five blocks along time. The next two Conv3D layers had a filter size of 1 along time, practically processing the 5 blocks separately. Finally, the extracted information was fused along time by the last Conv3D layer and by the topmost two fully connected layers.

As regards the STN, there are several options to extend it to 3D. First, we could reformulate the affine transformation to operate on 3D data. However, for our ultrasound data, translation, shearing and rotation along the time axis seemed to be unnecessary. Hence, we chose to keep the transformation in 2D. The second question was whether the localization network should operate on 2D or 3D data. We chose the technically simpler solution, and used the same 2D localization network as for the 2D-CNN. In this configuration, the STN applies the same 2D transformation to all the 25 images of the actual input block.

5. Results and discussion

First of all, we examined whether the inclusion of the STN module has any positive effect on the results when using only training data from one speaker and one session. Table 1 shows the MSE scores of the 2D network, with and without the STN mod-

Table 1: *MSE rates of the 2D-CNN on the dev set of the 4 speakers (top), and for the 4 extra sessions of speaker 048 (down).*

	s048	s049	s102	s103
no STN	0.361	0.387	0.390	0.343
with STN	0.358	0.396	0.389	0.340
	s048-2	s048-3	s048-4	s048-5
no STN	0.505	0.462	0.464	0.504
with STN	0.495	0.458	0.467	0.483

ule, for all the 4 speakers and the 4 extra sessions from speaker 048. In this configuration, we expected no significant benefit from the STN, and this is exactly what we see. The scores we obtained are quite similar across all speakers and the extra sessions of speaker 048. One notable difference is that the results are consistently worse for the additional sessions of speaker 048 than for her main session. The explanation is that these extra sessions were much shorter than the multi-speaker recordings (3.5 vs. 15 minutes).

In the adaptation experiments we considered the network trained on the data of speaker 048 as our base model. First, we evaluated it on the samples from the other speakers and the additional sessions of the same speaker without any adaptation. These results will be considered as the base results (i.e. 0% relative improvement due to adaptation). Then, we trained this baseline model further on the data from the other speakers and sessions, using three adaptation strategies. First, we allowed the full network to learn on the new data, and the results obtained this way will be considered to have 100% improvement (i.e. the best possible result achievable by adaptation). Next, we froze the weights of the regression network, and allowed only the STN module to adapt. In the final configuration, besides the STN we also allowed the re-adjustment of the weights in the output layer. This layer simply calculates a weighted linear combination of the hidden representations, so allowing it to adapt might significantly decrease the regression error, while it contains only about 1% of all the weights in the network. In all the adaptation experiments we used an initial learning rate that is an order of magnitude smaller than that of the baseline, but otherwise there were no other changes in the training process.

Table 2 summarizes the results achieved with the different adaptation strategies for the additional speaker (upper panel) and for the additional sessions of the same speaker (lower panel). As the relative improvements are more informative than the actual MSE values, in parentheses we display the relative error reductions (considering the no-adaptation case as 0% and full adaptation as 100%), and in the bottom line we summarize the average improvements. The first column ('None', i.e. no adaptation) clearly shows that the results are unacceptable without any adaptation – the error rates are actually worse than those for a randomly initialization net without any training. Moreover, the scores in the 'Full' adaptation column are pretty similar to those in the baseline table, so pre-training on speaker 048 seems to be neither beneficial nor detrimental (interestingly even for the extra sessions from the same speaker). By allowing only the STN module to adapt ('STN' column) we can eliminate 75-76% of the performance gap between the non-adapted and fully adapted models. This is pretty good, considering that the regression network itself is not modified at all, we just simply allow the STN to learn a more optimal affine transformation for the images of the given speaker or session. The improvement is even larger if we also allow the linear output layer of the regres-



Figure 3: A person has been recorded his speech and tongue movements by the Ultrasound Tongue Imaging device.

Table 2: MSE rates of the 2D-CNN model of spk048 for the other 3 speakers (top), and for the 4 extra sessions of the same speaker (down), using various adaptation strategies.

Speaker	Adaptation method				Mean θ
	None	STN	STN + out	Full	
spk049	1.049	0.588 (-71%)	0.517 (-82%)	0.400	0.887 (-25%)
spk102	1.401	0.609 (-78%)	0.449 (-94%)	0.389	1.015 (-38%)
spk103	1.322	0.552 (-79%)	0.469 (-88%)	0.350	0.909 (-42%)
Average diff.	-0%	-76%	-88%	-100%	-35%
Session	Adaptation method				Mean θ
	None	STN	STN + out	Full	
ses-2	1.131	0.646 (-77%)	0.547 (-93%)	0.503	0.913 (-34%)
ses-3	0.998	0.619 (-69%)	0.485 (-94%)	0.451	0.934 (-11%)
ses-4	1.054	0.641 (-70%)	0.522 (-90%)	0.468	0.908 (-24%)
ses-5	1.174	0.604 (-85%)	0.566 (-91%)	0.506	0.955 (-32%)
Average diff.	-0%	-75%	-92%	-100%	-26%

Table 3: MSE rates for cross-speaker adaptation, using the 3D-CNN network.

Speaker	Adaptation method			
	None	STN	STN + out	Full
spk049	1.105	0.553 (-73%)	0.497 (-80%)	0.348
spk102	1.451	0.502 (-84%)	0.416 (-91%)	0.315
spk103	1.541	0.501 (-83%)	0.418 (-90%)	0.294
Avg.diff.	-0%	-80%	-87%	-100%

sion network to adjust its weights to the new speaker or session ('STN + out'). In this case, the average error reduction is 88% for cross-speaker and 92% for cross-session adaptation. The better cross-session score is reasonable, as one would expect that the differences caused by the misalignment of the device might be easier to compensate by an affine transformation than the inherent anatomical differences between speakers.

As in the baseline evaluation (Table 1) the STN had no effect on the results, we speculated that maybe it is unnecessary to learn a separate transformation for *each image* – perhaps learning one global θ for *each speaker or session* would be enough. We examined the within-speaker variance of the 6 components of theta for the baseline models, and we indeed found that it was very low, on the order of 0.01, while the between-speaker variance of the θ vectors was 3-5 times larger. This fact seemed to underpin our conjecture, so we performed the following simple experiment. Instead of using a unique theta for each image, we simply replaced the STN of the adapted models with the mean θ vector over the samples of the given speaker/session. The results are shown in the rightmost columns of Table 2 ('mean θ '), but the findings are negative: while there is a 25-35% reduction in the error rates, it is very far from the best achievable. It seems that even though the variance of θ is small, the minor nuances in the learned transformations per image play an important role in the regression accuracy.

Finally, we extended our experiments to 3D input blocks consisting of 25 subsequent images, using a 3D-CNN for the regression task and a 2D-CNN for the STN module (cf. Section 4.3). In this case, the baseline model yielded an MSE of 0.275 and 0.278 with and without the STN on the data set of speaker 048. As speaker and session adaptation gave very simi-

lar results previously, we repeated only the cross-speaker adaptation experiments with this network, and the results are shown in Table 3. While all the error rates are typically lower than they were for the 2D-CNN, the relative improvements with respect to the various adaptation strategies are very similar.

6. Conclusions

Current tongue ultrasound-based SSI systems are sensitive to changing speakers, or even to a slight displacement of the recording device. In this study we examined whether an STN module is able to counterbalance these factors by applying an affine transform on the input image. When applying a simple 2D-CNN for spectral estimation, we found that allowing only the adaptation of the STN module can reduce the error rate by about 75%, while allowing also the linear output layer to adapt can compensate for 88-92% of the error. We also extended the experiments to 3D input blocks, and we observed similar tendencies, although the improvement was somewhat smaller. Considering, however, that the STN module and the output layer altogether contain only 10% of the weights of the full network, and that the improvements were consistent for cross-session and cross-speaker set-ups, the proposed method might allow for a quicker adaptation of the UTI-processing workflow. In the future we plan to run experiments with a 3D localization network and with smaller amounts of adaptation material.

7. Acknowledgements

The authors would like to express their gratitude to the following institutions and funding sources for their valuable support in conducting this research. This work was supported by the European Union project RRF-2.3.1-21-2022-00004 within the framework of the Artificial Intelligence National Laboratory and project TKP2021-NVA-09, implemented with the support provided by the Ministry of Innovation and Technology of Hungary from the National Research, Development, and Innovation Fund, financed under the TKP2021-NVA funding scheme and FK 142163 grant. T.G.Cs. was supported by the Bolyai János Research Fellowship of the Hungarian Academy of Sciences and by the ÚNKP-22-5-BME-316 grant. The RTX A5000 GPU card used was donated by the NVIDIA Corporation.

8. 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] 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, 2020.
- [3] 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.
- [4] A. Ephrat and S. Peleg, "Vid2speech: Speech reconstruction from silent video," in *Proc. ICASSP*, 2017, pp. 5095–5099.
- [5] 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," *arXiv preprint arXiv:2104.13332*, 2021.
- [6] D. Oneață, A. Stan, and H. Cucu, "Speaker disentanglement in video-to-speech conversion," in *Proc. EUSIPCO*, 2021, pp. 46–50.
- [7] N. Saleem, J. Gao, M. Irfan, E. Verdu, and J. P. Fuente, "E2E-V2SResNet: Deep residual convolutional neural networks for end-to-end video driven speech synthesis," *Image and Vision Computing*, vol. 119, p. 104389, 2022.
- [8] D. Wang, S. Yang, D. Su, X. Liu, D. Yu, and H. Meng, "VCVTS: Multi-speaker Video-to-Speech synthesis via cross-modal knowledge transfer from voice conversion," in *Proc. ICASSP*, 2022.
- [9] B. Denby and M. Stone, "Speech synthesis from real time ultrasound images of the tongue," in *Proc. ICASSP*, 2004, pp. 685–688.
- [10] 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*, 2017, pp. 3672–3676.
- [11] E. Tatulli and T. Hueber, "Feature extraction using multimodal convolutional neural networks for visual speech recognition," in *Proc. ICASSP*, 2017, pp. 2971–2975.
- [12] 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.
- [13] T. Csapó, M. Al-Radhi, G. Németh, G. Gosztolya, T. Grósz, L. Tóth, and A. Markó, "Ultrasound-based silent speech interface built on a continuous vocoder," 2019.
- [14] N. Kimura, M. Kono, and J. Rekimoto, "Sottovoce: An ultrasound imaging-based silent speech interaction using deep neural networks," in *Proc. of CHI Conf. on Human Factors in Computing Systems*, 2019.
- [15] T. G. Csapó, C. Zainkó, L. Tóth, G. Gosztolya, and A. Markó, "Ultrasound-based Articulatory-to-Acoustic Mapping with Wave-Glow Speech Synthesis," in *Proc. Interspeech*, 2020, pp. 2727–2731.
- [16] C. Zainkó, L. Tóth, A. H. Shandiz, G. Gosztolya, A. Markó, G. Németh, and T. G. Csapó, "Adaptation of Tacotron2-based Text-To-Speech for Articulatory-to-Acoustic Mapping using Ultrasound Tongue Imaging," in *Proc. ISCA SSW11*, 2021, pp. 54–59.
- [17] A. H. Shandiz, L. Tóth, G. Gosztolya, A. Markó, and T. G. Csapó, "Neural Speaker Embeddings for Ultrasound-Based Silent Speech Interfaces," in *Proc. Interspeech 2021*, 2021, pp. 1932–1936.
- [18] T. G. Csapó and K. Xu, "Quantification of Transducer Misalignment in Ultrasound Tongue Imaging," in *Proc. Interspeech*, 2020, pp. 3735–3739.
- [19] 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.
- [20] 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 *Proc. SLT*, 2021, pp. 1109–1116.
- [21] J.-X. Zhang, K. Richmond, Zhen-Hua-Ling, and L.-R. Dai, "TaL-Net: Voice Reconstruction from Tongue and Lip Articulation with Transfer Learning from Text-to-Speech Synthesis," in *Proc. AAAI*, 2021.
- [22] Y. Yu, A. H. Shandiz, and L. Tóth, "Reconstructing speech from real-time articulatory mri using neural vocoders," in *Proc. EUSIPCO*, 2021, pp. 945–949.
- [23] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [24] S. Ji, W. Xu, M. Yang, and K. Yu, "3d convolutional neural networks for human action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 221–231, 2013.
- [25] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri, "A closer look at spatiotemporal convolutions for action recognition," in *Proc. CVPR*, 2018.
- [26] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6202–6211.
- [27] L. Tóth and A. H. Shandiz, "3D convolutional neural networks for ultrasound-based silent speech interfaces," in *Proc. ICAISC*. Springer, 2020, pp. 159–169.
- [28] P. Saha, Y. Liu, B. Gick, and S. Fels, "Ultra2speech – a deep learning framework for formant frequency estimation and tracking from ultrasound tongue images," in *Proc. MICCAI*, 2020, pp. 473–482.
- [29] G. Bertasius, H. Wang, and L. Torresani, "Is space-time attention all you need for video understanding," *Proc. ICML*, vol. 139, pp. 813–824, 2021.
- [30] D. Serdyuk, O. Braga, and O. Siohan, "Transformer-based video front-ends for audio-visual speech recognition," *arXiv preprint arXiv:2201.10439*, 2022.
- [31] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," in *NIPS* 28, 2015, pp. 2017–2025.
- [32] A. Bas, P. Huber, W. A. P. Smith, M. Awais, and J. Kittler, "3d morphable models as spatial transformer networks," in *Proc. Int. Conf. on Computer Vision (ICCV) Workshops*, 2017, pp. 895–903.
- [33] J. Yu and S. Wang, "Visual speech recognition in natural scenes based on spatial transformer networks," in *Proc. Int. Conf. on Anti-counterfeiting, Security, and Identification (ASID)*, 2020.
- [34] W. Chu, W.-C. Hung, Y.-H. Tsai, D. Cai, and M.-H. Yang, "Weakly-supervised caricature face parsing through domain adaptation," in *Proc. Int. Conf. on Image Processing (ICIP)*, 2019, pp. 3282–3286.
- [35] R. Robinson, Q. Dou, D. Coelho de Castro, K. Kamnitsas, M. de Groot, R. M. Summers, D. Rueckert, and B. Glocker, "Image-level harmonization of multi-site data using image-and-spatial transformer networks," in *Proc. MICCAI*, 2020, pp. 710–719.
- [36] C. Zhao, J. Zhang, C. Wu, H. Wang, and K. Xu, "Predicting tongue motion in unlabeled ultrasound video using convolutional LSTM neural networks," in *Proc. ICASSP*, 2019, pp. 5926–5930.