Improved Processing of Ultrasound Tongue Videos by Combining ConvLSTM and 3D Convolutional Networks

Amin Honarmandi Shandiz, László Tóth

University of Szeged, Institute of Informatics, Hungary
{shandiz,tothl}@inf.u-szeged.hu

Abstract. Silent Speech Interfaces aim to reconstruct the acoustic signal from a sequence of ultrasound tongue images that records the articulatory movement. The extraction of information about the tongue movement requires us to efficiently process the whole sequence of images, not just as a single image. Several approaches have been suggested to process such a sequential image data. The classic neural network structure combines two-dimensional convolutional (2D-CNN) layers that process the images separately with recurrent layers (eg. an LSTM) on top of them to fuse the information along time. More recently, it was shown that one may also apply a 3D-CNN network that can extract information along both the spatial and the temporal axes in parallel, achieving a similar accuracy while being less time consuming. A third option is to apply the less well-known ConvLSTM layer type, which combines the advantages of LSTM and CNN layers by replacing matrix multiplication with the convolution operation. In this paper, we experimentally compared various combinations of the above mentions layer types for a silent speech interface task, and we obtained the best result with a hybrid model that consists of a combination of 3D-CNN and ConvLSTM layers. This hybrid network is slightly faster, smaller and more accurate than our previous 3D-CNN model.

Keywords: Silent Speech Interface, Convolutional neural network, 3D convolution, ConvLSTM, Ultrasound Tongue Video

1 Introduction

The area of Silent Speech Interfaces (SSI) deals with the problem of converting articulatory recordings to speech signals [1]. The studies in this field have considered various types of articulatory signals as input, such as Electroencephalography (EEG), Electromagnetic Articulography (EMA), Ultrasound Tongue Video Imaging (UTI) and so on. SSIs could provide a great amount of help for those disabled people who can not talk loud, but are able to silently articulate the speech. Converting the signals recorded during articulation to speech would allow these people to interact with others. SSI solutions could also be used in some other conditions where normal communication is not possible, for example in certain

military situations or in very noisy industrial environments where people could barely understand each other. In this paper, we work with Ultrasound Tongue Video Imaging (UTI) Data [2,3,4] recorded from Hungarian speakers.

In SSI systems, the classic approach of estimating the speech signal from the articulatory data consisted of two steps: the first one is to to convert the input to text using speech recognition, and the second is to synthesize speech based on the text. Nowadays, however, directly converting the articulatory signals to speech is more popular, as it is less time consuming and seems to be more suitable for real-time applications. This direct approach has been made viable by the Deep Neural Network (DNN) technology, which also revolutionized many other speech-related fields, for example speech recognition [5], and speech synthesis [6]. Here, we also follow the direct approach and apply DNNs [7,3,8].

The given articulatory-to-acoustic mapping task can be addressed by applying simple fully connected DNNs [7,4]. However, as we are working with images, using convolutional neural networks (CNN) [9] is more reasonable, and they have been successfully applied to the SSI task by several authors [10,3]. A very important further aspect is that that our input consists of a sequence of images (an ultrasound video), so it is not effective to simply process these images separately. Thus, we can apply a Recurrent Neural Network (RNN) such as an LSTM [11] to extract information from a sequence. And when the sequence consists of images, like in our case, an obvious solution is to combine an LSTM with a 2D CNN that extracts the information from individual video frames [12]. Alternatively, it would be possible to extend the 2D convolution to 3D by adding the time axis as an extra dimension [13]. This approach was followed in [3], and it was found that the two approaches can achieve very similar performances. However, there is a third option: for the processing of image sequences, Shi et al proposed the ConvLSTM layer type [14], which combines the advantages of convolutional and recurrent processing in one layer. For some reason, however, the ConvLSTM construct is not widely known. To knowledge, it was applied to UTI data only in one case [15], but even in that paper the task was different from ours. The goal of this paper is to experiment with ConvLSTM models, and compare their performance with the previous 2D-CNN+LSTM and 3D-CNN approaches. We also try to combine the three types of layers, resulting in hybrid models, and our results show that the hybrid approach yields the best performance for our task.

The paper is organized as follows. In Section 2, we briefly introduce the concept of the Convolutional LSTM that we are going to use. In Section 3, we explain the data acquisition and processing steps for our input and output data. In Section 4, we present our experimental setup, while in and following Section 5, the experimental results are discussed and explained. Finally, in Section 6 our main conclusions are given.

2   Convolutional LSTM for SSI

SSI systems synthesize speech from articulatory videos by learning the mapping between the input ultrasound image sequence and the output audio signal. SSI is
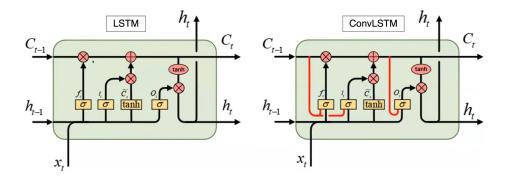
Fig. 1. *Internal structure of a standard LSTM cell and its extended version (with extra peephole connections) used in Convolutional LSTMs [17,18].*

LSTM

$$i_t = \sigma(W_{xi}x_t + W_{hi}h_{t-1} + W_{ci} \circ c_{t-1} + b_i)$$
$$f_t = \sigma(W_{xf}x_t + W_{hf}h_{t-1} + W_{cf} \circ c_{t-1} + b_f)$$
$$c_t = f_t \circ c_{t-1} + i_t \circ \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c)$$
$$o_t = \sigma(W_{xo}x_t + W_{ho}h_{t-1} + W_{co} \circ c_t + b_o)$$
$$h_t = o_t \circ \tanh(c_t)$$

ConvLSTM

$$i_t = \sigma(W_{xi} * \mathcal{X}_t + W_{hi} * \mathcal{H}_{t-1} + W_{ci} \circ \mathcal{C}_{t-1} + b_i)$$
$$f_t = \sigma(W_{xf} * \mathcal{X}_t + W_{hf} * \mathcal{H}_{t-1} + W_{cf} \circ \mathcal{C}_{t-1} + b_f)$$
$$\mathcal{C}_t = f_t \circ \mathcal{C}_{t-1} + i_t \circ \tanh(W_{xc} * \mathcal{X}_t + W_{hc} * \mathcal{H}_{t-1} + b_c)$$
$$o_t = \sigma(W_{xo} * \mathcal{X}_t + W_{ho} * \mathcal{H}_{t-1} + W_{co} \circ \mathcal{C}_t + b_o)$$
$$\mathcal{H}_t = o_t \circ \tanh(\mathcal{C}_t)$$

Fig. 2. *The equations behind the operation of Long Short Term Memory (LSTM) versus Convolutional LSTM neurons [14].*

a sequential task, as both the input and the output are sequences, with a strong correlation between consecutive elements of the sequence. As in our case the input data consist of ultrasound images, convolutional networks (CNN) seems to be a proper tool for processing the input, as they are known to perform well when working with images [9], and also in particular with SSI ultrasound tongue images [16,12]. As our input is a sequence, the information content along the time axis of the data can be extracted by applying Recurrent Neural Networks (RNNs). In particular, a variant of recurrent networks called the Long-Short Term Memory (LSTM) is known to be more effective in extracting long-term dependencies in the input sequence [11]. These networks have special gates in their internal implementations which improve their abilities to handle large-distance relations between time-related features.

The data flow in the standard implementation of an LSTM is shown in Fig. 1. Some implementations also contain extra connections, such a so-called "peephole" variant in shown on the right side of Fig. 1, In both cases, the input of the LSTM consist of a sequence of vectors. As in our case we want to process a sequence of images, a straightforward solution is to combine a CNN with an LSTM. In the trivial arrangement the images of the input are first processed by a (2-dimensional) CNN, and the sequence of CNN outputs are integrated over time by using an LSTM. This approach was shown to work fine in SSI imple-

mentations [12]. and we will shortly refer to this scheme as the CNN+LSTM approach. However, this type of processing requires the combination of two layers – a 2D-CNN to process the data along the two spacial axes, and an LSTM to process it along the temporal axis. A more efficient solution called the Convolutional LSTM, or shortly ConvLSTM has been proposed by Shi et al. As the name says, their solution performs the two processing steps in one. It can extract spatio-temporal features from the input data by applying a convolution operation in the inner steps LSTM instead of matrix multiplication [14]. This is reflected in the equations of Fig. 2. where "$*$" represents the convolution operation, and "$\circ$" stands for gating. Note that, apart from the convolution, the equations are exactly the same as those of the (peephole) LSTM. The convolution allows the more efficient processing of image sequences, resulting in better performance with much fewer parameters. For example, Kwon et al. successfully applied a hierarchical ConvLSTM for speech recognition [19]. Recently, Zhao et al. used ConvLSTMs for predicting subsequent ultrasound images in an SSI task [15].

Processing a sequence of images is also viable by extending the convolution operation to the time axis, resulting in a three-dimensional convolution (3D-CNN) model. The main advantage of this approach is that it is faster, as it applies only convolution operations. Convolution also allows the skipping of input images, which is not possible in an LSTM framework. In a previous paper, the 3D-CNN model gave results that were comparable or slightly better than those with the more conventional CNN+LSTM approach [3]. Here, we extend this earlier comparison to the ConvLSTM model, and we are also going to experiment with hybrid models that combine 3D-CNN and ConvLSTM layers.

3   Data acquisition and preprocessing

The ultrasound data was collected from a Hungarian female subject (42 years old) while she was reading sentences aloud. Her tongue movement was recorded in a midsaggital orientation – placing the ultrasonic imaging probe under the jaw – using a "Micro" ultrasound system by Articulate Instruments Ltd. The transducer was fixed using a stabilizer headset. The 2-4 Mhz / 64 element 20 mm radius convex ultrasound transducer produced 82 images per second. The speech signals were recorded in parallel with an Audio-Technica ATR 3350 omnidirectional condenser microphone placed at a distance of 20 cm from the lips. The ultrasound and the audio signals were synchronized using the software tool provided with the equipment. Altogether 438 sentences (approximately half an hour) were recorded from the subject, which was divided into train, development and test sets in a 310-41-87 ratio. We should add that the same data set was used in several earlier studies [2,7,3], and the data set is publicly available[1]. The ultrasound probe records 946 samples along each of its 64 scan lines. The recorded data can be converted to conventional ultrasound images using the software tools provided. However, due to its irregular shape, this image is harder

---

[1] The dataset is available upon request from csapot@tmit.bme.hu

to process by computers, while it contains no extra information compared to the original scan data. Hence, we worked with the original 964x64 data items, which were downsampled to 128x64 pixels. The intensity range of the data was min-max normalized to the [-1, 1] interval before feeding it to the network.

The speech signal was recorded with a sampling rate of 11025 Hz, and then converted to a 80-bin mel-spectrogram using the SPTK toolkit (http://sptk.sourceforge.net). The goal of the machine learning step was to learn the mapping between the sequence of ultrasound images and the sequence of mel-spectrogram vectors. As the two sequences are perfectly synchronized, it was not necessary to apply a sequence-to-sequence learning strategy. We simply defined the goal of learning as an image-to-vector mapping task, using the mean squared error (MSE) as the loss function in the network training step. The 80 mel-frequency coefficients served as training targets, from which the speech signal was reconstructed using WaveGlow [20]. To facilitate training, each of the 80 targets were standardized to zero mean and unit variance. The input of training consisted of a block of 25 consecutive. This allowed all DNN variants to involve the time axis in the information extraction process. The whole SSI framework followed our earlier study [3].

4   Experimental Setup

We implemented our networks using Keras with a tensorflow back-end [21]. We applied three different network architectures that can process 3-dimensional blocks of data. In the tables, "3D-CNN" refers to the fully convolutional model proposed in [3]. This model does not have any LSTM component. "3D-CNN + BiLSTM" refers to a combination which applies a BiLSTM layer as the topmost hidden layer to integrate the temporal features extracted by the previous 3D-CNN layers. The final model referred to as "3D-CNN + ConvLSTM" replaces the tompost 3D-CNN and BiLSTM layers by a ConvLSTM layer. Notice that the ConvLSTM technique fuses the convolution and the LSTM operations into one layer, so here we can also spare one hidden layer by this substitution. In the following, we give a more detailed description of the three configurations.

3D Convolutional Neural Network(3D-CNN): This model was described in detail in [3], and its network layers are shown in Table 1. The networks processes the input sequence of 25 video frames in 5-frame blocks using 3D convolution. The overlap between these blocks is minimized by setting the stride parameter s of the time axis to 5. These blocks are processed further by 3 additional Conv3D layers, with pooling layers after every second convolution layer. Finally, the output is flattened and integrated over the time axis by a dense layer as the topmost hidden layer. The output hidden layer is a linear layer with 80 neurons, corresponding to the 80 spectral parameters given as training targets. This special network structure was motivated by Tran et al., who found that for the best result the processing should focus on the two spacial axes first, performing the integration over the temporal axis only afterwards [13]. Toth at al. also obtained the best result with performing the 3D convolution in this decomposed,

"(2D+1)" form [3]. Compared to that study, we achieved slightly better results with the same architecture by switching to the Adam optimizer instead of SGD, and by adjusting some meta-parameters, for example the dropout rate.

3D CNN + BiLSTM: As the output of the four layers of 3D convolution, the 3D-CNN network produces a sequence of 5 matrices, which are combined by a simple dense layer (cf. Table 1). Our first modification was to replace this fully connected layer with a (bidirectional) LSTM, which required us to reshape the matrices into vectors. The LSTM is a more sophisticated solution to extract the information from a temporal sequence, so we hoped to get slightly better results from this approach. As Table 1 shows, we set the *return_sequences* parameter of the LSTM to False, so the output is a simple vector, which serves as the input of the subsequent dense linear output layer.

3D CNN + ConvLSTM: Our main goal in this paper was to examine the efficiency of the ConvLSTM layer for this task. In the first experiment we applied it only at the topmost hidden layer of our 3D-CNN model (see Table 1). As the ConvLSTM layer implements the operation of a convolutional and an LSTM layer in one, we replaced the uppermost Conv3D and LSTM layers by a ConvLSTM layer, reducing the number of neural hidden layers from 5 to 4. Also, as the ConvLSTM layer works with matrices and also outputs matrices, the reshaping was required after the layer and not before it.

5   Results and Discussion

In the first experiment we compared the performance of the baseline 3D-CNN model we the two hybrid solutions proposed in the previous chapter and in Table 1. In Table 2 we report two simple objective metrics of the quality of training, the mean squared error (MSE) and the R^2 score, which is popular in regression tasks implemented with neural networks (for R^2 a higher value means better performance). As can be seen, replacing the dense layer by the LSTM layer already brings a slight but consistent improvement in the results, both on the development and on the test set. Fusing the uppermost Conv3D and the LSTM layer into a ConvLSTM layer resulted in further error reduction of about the same rate, even though the network depth is decreased. This clearly proves the efficiency of the ConvLSTM layer. However, we also observed a drawback, namely that the ConvSLTM layer has much more trainable parameters than the Conv3D layer. Hence, we had to reduce the filter size in the ConvLSTM layer, in order to keep the number of parameters in the original range. Theoretically, similar to the LSTM layer, the ConvLSTM layer can also be made bidirectional. However, we ran into the same problem that it tremendously increased the number of parameters while yielding only a marginal improvement. Thus, we stuck with using the unidirectional variant. Finally, to fuse the Conv3D and the LSTM layers, we had to remove the second MaxPooling layer. We also tried to insert it back after the ConvLSTM layer, but the results did not change considerably.

Obviously, many other possible configurations exist that combine Conv3D and ConvLSTM layers. In the second experiment we tried out further combina-

Table 1. The layers of the 3D-CNN, the 3D-CNN + BiLSTM and the 3D-CNN + ConvLSTM networks in Keras implementation, along with their most important parameters. The differences are highlighted in bold.

3D-CNN	3D-CNN + BiLSTM
Conv3D(30, (5,13,13), strides=(s, 2,2)) Dropout(0.3)	Conv3D(30,(5,13,13), strides=(s,2,2)) Dropout(0.3)
Conv3D(60, (1,13,13), strides=(1,2,2)) Dropout(0.3) MaxPooling3D(pool_size=(1,2,2))	Conv3D(60,(1,13,13),strides=(1,2,2)) Dropout(0.3) MaxPooling3D(pool_size=(1,2,2))
Conv3D(90, (1,13,13), strides=(1,2,1)) Dropout(0.3)	Conv3D(90,(1,13,13),strides=(1,2,1)) Dropout(0.3)
Conv3D(85, (1,13,13), strides=(1,2,2)) Dropout(0.3) MaxPooling3D(pool_size=(1,2,2))	Conv3D(85, (1,13,13), strides=(1,2,2)) Dropout(0.3) MaxPooling3D(pool_size=(1,2,2))
Flatten() **Dense(500)** Dropout(0.3)	**Reshape((5, 340))** **Bidirectional(LSTM(320,** **ret_seq=False))**
Dense(80, activation='linear')	Dense(80, activation='linear')

3D-CNN + ConvLSTM
Conv3D(30,(5,13,13),strides=(s,2,2)) Dropout(0.35)
Conv3D(60,(1,13,13),strides=(1,2,2)) Dropout(0.35) MaxPooling3D(poolsize =(1,2,1))
Conv3D(90,(1,13,13),strides=(1,2,2)) Dropout(0.35) **ConvLSTM2D(64, (3,3), Strides=(2,2), ret_seq=False)** **Flatten()**
Dense(80,activation='linear')

tions of these two layers. We experimented with 4 hidden layer constructs, and we fixed the uppermost layer to be a ConvLSTM, as it convincingly proved to be the better setup in the previous experiment. Table 3 summarizes the architectures we experimented with. As regards ConvLSTM layers, the *return_sequences* parameter was set to True for intermediate layers, and set to False only when the ConvLSTM layer was the topmost hidden layer. The meta-parameters were always chosen so that the global count of the free parameters stayed similar to that of the baseline model.

Seeing the good performance of the ConvLSTM layer in the previous experiment, we first tried to build a fully ConvLSTM model. However, as the first row of Table 3 shows, we obtained no improvement. As the ConvLSTM layer proved to be more efficient than the Conv3D layer earlier, next we tried to create a

Table 2. The MSE and mean R^2 scores obtained with the various network configurations for the development and test sets, respectively. The best results are highlighted in bold.

Network Type	Dev		Test	
	MSE	Mean R^2	MSE	Mean R^2
3D-CNN	0.292	0.714	0.293	0.710
3D-CNN + BiLSTM	0.285	0.721	0.282	0.721
3D-CNN + ConvLSTM	**0.276**	**0.727**	**0.276**	**0.73**

Table 3. The MSE for different combinations of Conv3D and ConvLSTM layers in the four hidden layers of the network. The best results are highlighted in bold.

Layer 1	Layer 2	Layer 3	Layer 4	Dev	Test
ConvLSTM	ConvLSTM	ConvLSTM	ConvLSTM	0.31	0.31
ConvLSTM	ConvLSTM	ConvLSTM	—	0.29	0.3
Conv3D	ConvLSTM	ConvLSTM	ConvLSTM	0.31	0.31
Conv3D	Conv3D	ConvLSTM	ConvLSTM	0.36	0.35
Conv3D	**Conv3D**	**Conv3D**	**ConvLSTM**	**0.27**	**0.27**
Conv3D	ConvLSTM	Conv3D	ConvLSTM	0.3	0.3
ConvLSTM	Conv3D	Conv3D	ConvLSTM	0.34	0.34

network of just 3 ConvLSTM layers instead of 4. As shown in the second row of the table, the results became slightly better, but still worse than the baseline. This result reinforces our previous observation that ConvLSTM networks do not require the same depth as a convolutional network.

Conv3D layers have the advantage that they can easily downsample the time axis using a stride parameter larger than 1. On the contrary, ConvLSTM units cannot easily skip elements of their input sequence, due to their recurrent nature, which results in a large parameter count and a slow training. Hence, it seemed to be more efficient to put a Conv3D layer into the first hidden layer. We tried to place 1-2-3 Conv3D layers in the lower layers, and 3-2-1 ConvLSTM layers in the remaining layers. The middle block of Table 3 shows that the optimal solution is to have just one ConvLSTM layer, as in our original experiment. Lastly, we tried two further configurations with alternating Conv3D and ConvLSTM layers, motivated by papers like [19,15], but we did not receive any better results.

6   Conclusion

Here, we were seeking the optimal neural network architecture for the articulatory-to-acoustic mapping task of SSI systems. The task involves the processing of 3D data blocks – sequences of images – for which one can apply 3D-CNN models,

such as in [3]. Alternatively, one may apply a ConvLSTM model proposed by [14]. Besides comparing the purely convolutional and ConvLSTM models, we also experimented with hybrid architectures where the two layers types are mixed. The 3D-CNN + ConvLSTM hybrid model obtained the best results, better than the baseline 3D-CNN model, and it also outperformed other models with a different order of layers, as applied in [19] for emotion recognition, and in [15] for the prediction of the subsequent ultrasound image. Applying the ConvLSTM layer in the uppermost hidden layer even made the model smaller (with one hidden layer) and slightly faster to train. The optimal model arrangement consists of three Conv3D layers and a ConvLSTM on top of them, which illustrates that it is worth combining the ConvLSTM layer with other layer types such as the Conv3D to extract spatio-temporal features from videos – in our case, to better capture the tongue movement. The winning architecture also shows that the Conv3D blocks are more efficient in extracting local spectro-temporal information, while ConvLSTM is more efficient in fusing these pieces of information along the time axis. Interestingly, this coincides with the observation of Tran et al. about the optimal order of feature extraction for 3D video blocks [13]. In the future we plan to extend our research to transformer models that apply two separate networks for encoding and decoding, which would allow us to experiment with different network types for the decoder and encoder components, similar to the UNET [22] architecture.

7 Acknowledgments

References

1. Schultz, T., Wand, M., Hueber, T., Krusienski, D.J., Herff, C., Brumberg, J.S.: Biosignal-based spoken communication: A survey. IEEE/ACM Transactions on Audio, Speech, and Language Processing **25**(12), 2257–2271 (2017)
2. Csapó, T.G., Grósz, T., Gosztolya, G., Tóth, L., Markó, A.: DNN-based ultrasound-to-speech conversion for a silent speech interface. In: Proc. Interspeech. pp. 3672–3676 (2017)
3. Tóth, L., Shandiz, A.H.: 3d convolutional neural networks for ultrasound-based silent speech interfaces. In: International Conference on Artificial Intelligence and Soft Computing. pp. 159–169. Springer (2020)
4. Jaumard-Hakoun, A., Xu, K., Leboullenger, C., Roussel-Ragot, P., Denby, B.: An articulatory-based singing voice synthesis using tongue and lips imaging. In: ISCA Interspeech 2016. vol. 2016, pp. 1467–1471 (2016)

5. Hinton, G., Deng, L., Yu, D., Dahl, G.E., Mohamed, A.r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T.N., et al.: Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. IEEE Signal processing magazine **29**(6), 82–97 (2012)
6. Ling, Z.H., Kang, S.Y., Zen, H., Senior, A., Schuster, M., Qian, X.J., Meng, H.M., Deng, L.: Deep learning for acoustic modeling in parametric speech generation: A systematic review of existing techniques and future trends. IEEE Signal Processing Magazine **32**(3), 35–52 (2015)
7. Grósz, T., Gosztolya, G., Tóth, L., Csapó, T.G., Markó, A.: F0 estimation for dnn-based ultrasound silent speech interfaces. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 291–295. IEEE (2018)
8. Young, T., Hazarika, D., Poria, S., Cambria, E.: Recent trends in deep learning based natural language processing. ieee Computational intelligenCe magazine **13**(3), 55–75 (2018)
9. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems **25**, 1097–1105 (2012)
10. Saha, P., Liu, Y., Gick, B., Fels, S.: Ultra2speech-a deep learning framework for formant frequency estimation and tracking from ultrasound tongue images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 473–482. Springer (2020)
11. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
12. Juanpere, E.M., Csapó, T.G.: Ultrasound-based silent speech interface using convolutional and recurrent neural networks. Acta Acustica united with Acustica **105**(4), 587–590 (2019)
13. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proc. CVPR (2018)
14. Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. arXiv preprint arXiv:1506.04214 (2015)
15. Zhao, C., Zhang, P., Zhu, J., Wu, C., Wang, H., Xu, K.: Predicting tongue motion in unlabeled ultrasound videos using convolutional LSTM neural networks. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 5926–5930. IEEE (2019)
16. Kimura, N., Kono, M., Rekimoto, J.: Sottovoce: an ultrasound imaging-based silent speech interaction using deep neural networks. In: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. pp. 1–11 (2019)
17. Convolutional LSTM (2019), https://medium.com/neuronio/an-introduction-to-convlstm-55c9025563a7
18. Recurrent neural networks and LSTMs with keras (2020), https://blog.eduonix.com/artificial-intelligence/recurrent-neural-networks-lstms-keras
19. Kwon, S., et al.: CLSTM: Deep feature-based speech emotion recognition using the hierarchical ConvLSTM network. Mathematics **8**(12), 2133 (2020)
20. Prenger, R., Valle, R., Catanzaro, B.: Waveglow: A flow-based generative network for speech synthesis. In: Proc. ICASSP. pp. 3617–3621 (2019)
21. Chollet, F., et al.: Keras. https://github.com/fchollet/keras (2015)
22. Behboodi, B., Rivaz, H.: Ultrasound segmentation using u-net: learning from simulated data and testing on real data. In: 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). pp. 6628–6631. IEEE (2019)