

PREDICTING TONGUE MOTION IN UNLABELED ULTRASOUND VIDEOS USING CONVOLUTIONAL LSTM NEURAL NETWORKS

Chaojie Zhao¹, Peng Zhang², Jian Zhu³, Chengrui Wu⁴, Huaimin Wang⁵, Kele Xu^{5*}

¹Graduate School at Shenzhen, Tsinghua University, Shenzhen, China

²University of Science and Technology of China, Suzhou, China

³Department of Linguistics, University of Michigan, Ann Arbor, USA

⁴Columbia University in the City of New York, New York, USA

⁵National University of Defense Technology, Changsha, China

kelele.xu@gmail.com

ABSTRACT

A challenge in speech production research is to predict future tongue movements based on a short period of past tongue movements. This study tackles speaker-dependent tongue motion prediction problem in unlabeled ultrasound videos with convolutional long short-term memory (ConvLSTM) networks. The model has been tested on two different ultrasound corpora. ConvLSTM outperforms 3-dimensional convolutional neural network (3DCNN) in predicting the 9th frames based on 8 preceding frames, and also demonstrates good capacity to predict only the tongue contours in future frames. Further tests reveal that ConvLSTM can also learn to predict tongue movements in more distant frames beyond the immediately following frames. Our codes are available at: <https://github.com/shuiliwanwu/ConvLstm-ultrasound-videos>.

Index Terms— convolutional recurrent neural network, motion prediction, speech production, ultrasound tongue imaging, silent speech interface

1. INTRODUCTION

Speech articulation is characterized by coupled, transient and highly variable movements of various articulators. Despite the lack of invariance, tongue movements during speech can be predictable due to speaker-specific strategies and constraints imposed by the physical properties of the tongue and phonological structure of the language [1, 2, 3]. Understanding the dynamics of tongue motions not only is the central theoretical quest in speech production research but also inspires many potential practical applications, such as articulatory synthesis, silent speech interfaces and pronunciation training interfaces.

Tongue movements during speech production can be captured by the ultrasound imaging technique. Ultrasound

tongue imaging has been widely used in speech production and clinical research, as it provides a fast and non-invasive means to visualize the real-time movements of the tongue [4]. Extracting meaningful articulatory information from ultrasound image sequences remains a challenging task, though various methods have been proposed for contour tracking (e.g., [5]) and whole image analysis (e.g., [6, 7]). However, these methods are often limited to individual frame analysis of static tongue shapes. In this article, built on previous works on tongue motion prediction, we explore the use of convolutional long short-term memory network (ConvLSTM) in predicting tongue motions with unlabeled ultrasound videos.

2. RELATED WORK

Predicting tongue motions in ultrasound images remains largely unexplored in speech production research. Previous studies have primarily focused on extracting static tongue shape features from a single ultrasound frame (e.g., [6, 7, 8]). A recent work utilized 3-dimensional convolutional neural networks (3DCNN) to predict tongue motions in ultrasound videos, showing that future ultrasound frames can be predicted based on past frames [9]. In contrast to traditional convolutional neural networks, 3DCNNs explicitly incorporate the time dimension into the model by convolving the stacked consecutive input frames with 3-dimensional kernels [10] and therefore are suitable for processing ultrasound videos.

Apart from 3DCNN, long short-term memory (LSTM) structure has also been shown to be suitable for modeling long term dynamics of human motions and has been applied to human motion classification and prediction (e.g., [11, 12]). ConvLSTM has been proposed to specifically solve spatial-temporal sequence prediction problems [13]. In a ConvLSTM, the convolutional modules can capture the vocal tract features of each ultrasound frame at fine granularity, while LSTM cells exploit the auto-correlated properties of tongue

*Corresponding author.

movements to infer future frames. The present article describes the first use of ConvLSTMs to predict tongue motions in unlabeled ultrasound data with good performance.

The contributions of the current work can be summarized as follows: 1) We applied ConvLSTM to predict tongue motions in ultrasound image sequences based on past motions, demonstrating that in most cases ConvLSTM network outperforms the 3DCNN adopted in [9]; 2) [9] only investigated predictions of tongue motions in the immediately following frames. However, we also tested the performance of our method in predicting tongue motions in more distant ultrasound frames and ConvLSTM consistently outperforms 3DCNN in multiple quantitative comparisons.

3. CONVOLUTIONAL LSTM NEURAL NETWORK

The ConvLSTM is specifically designed for spatial-temporal sequence prediction problems [13]. The convolutional layers can effectively encode the salient visual features of the input images, while LSTM cells can model the temporal dependence across continuous tongue movements with better optimization stability. The structure of a convolutional LSTM cell is illustrated in Fig. 1. We adopted a slightly revised version of ConvLSTM [14], which can be formulated as follows:

$$i_t = \sigma(W_{xi} * X_t + W_{hi} * h_{t-1} + b_i) \quad (1)$$

$$f_t = \sigma(W_{xf} * X_t + W_{hf} * h_{t-1} + b_f) \quad (2)$$

$$\tilde{c}_t = \tanh(W_{x\tilde{c}} * X_t + W_{h\tilde{c}} * h_{t-1} + b_{\tilde{c}}) \quad (3)$$

$$c_t = f_t \circ c_{t-1} + i_t \circ \tilde{c}_t \quad (4)$$

$$o_t = \sigma(W_{xo} * X_t + W_{ho} * h_{t-1} + b_o) \quad (5)$$

$$h_t = o_t \circ \tanh(c_t) \quad (6)$$

where X_t represents the input tensor, h_t the hidden states, and c_t the cell states. And i_t , f_t and o_t are the gate activations at current time step. The “ \circ ” denotes the Hadamard product operation and “ $*$ ” denotes convolution operation. The weights (W) and biases (b) are learned through back propagation during training. In this study, the proposed model was built by stacking multiple ConvLSTM layers to model sequences of ultrasound frames, as illustrated in Fig. 2.

4. DATASET

Our experiments were based on the same two ultrasound corpora as in [9]. These ultrasound data were recorded as 320x240 pixel mid-sagittal ultrasound tongue images using an acquisition helmet and an ultrasound probe beneath the speaker’s chin. Detailed descriptions can be found in [9]. In all experiments, ultrasound images were resized to 96x96 pixels before further processing.

The “WSJ0” data were extracted from the Silent Speech Challenge corpus [8], which contains ultrasound video data

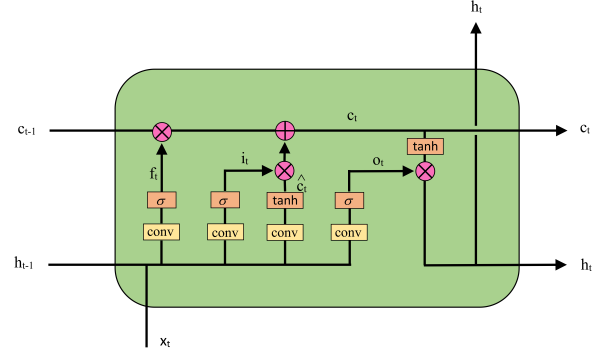


Fig. 1. Structure of a convolutional LSTM cell

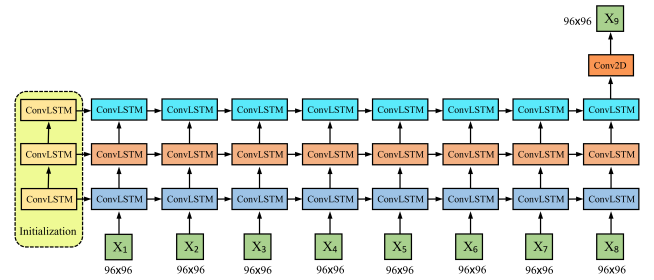


Fig. 2. A schematic illustration of the ConvLSTM network

at 60 Hz from a male native speaker producing non-vocalized sentences. The total number of training images were more than 700,000 frames of TIMIT sentences, while the number of test images were derived from 35,000 WSJ0 sentences.

“TJU” ultrasound data were collected at Tianjin University by recording a non-native female speaker reading a simple training passage (9900 images) and test passage (4800 images) in English. Ultrasound data were recorded at a frame rate of 30 Hz. A region of interest (ROI) containing the tongue was selected before resizing.

Tongue surface contours were extracted automatically with the classic snake algorithm [15] from the TJU dataset. Tongue extraction was only done on the TJU data because the tongue visibility in this data is better than that in the WSJ0 data. These contours represent the position of actual tongue surface in these ultrasound images. A sample can be found in the 3rd image of Fig. 4.

5. EXPERIMENTS

A series of experiments were carried out to test the performance of the ConvLSTM architecture in the tongue motion prediction task. The architecture developed is illustrated in Fig.3. It consists of three ConvLSTM layers, and a 2-dimensional convolution layer as the output layer. For all experiments, the inputs were of size 8×96×96, and the out-

puts of size $1 \times 96 \times 96$. We adopted the Adam optimizer with a learning rate of 0.001 and a batch size of 16.




Fig. 3. Structure of the ConvLSTM adopted in this study

- We first compared the performance of ConvLSTM with results in [9] using the same setting. For WSJ0 and TJU data respectively, the input data were generated by stacking 8 temporally consecutive frames together, and the outputs were the immediately following 9th frames. All images from TJU dataset were used in training but only a subset of images from WSJ0 were randomly selected for training due to limited computational resources.
- For the TJU data, as the snake-based contours were readily available, we also used the 8 consecutive frames to predict the snake-extracted contours in the 9th frame, which was the same as the “Cross” condition in [9]. In other words, the model is trained to predict the extracted tongue surface contours in the 9th frame.
- We also tested the model performance on predicting more distant tongue movements. In these experiments, 8 stacked consecutive ultrasound frames were used to predict the actual tongue image in the following 10th or the 11th frame, that is, future 2 or 3 frames immediately after the input frames. The prediction was evaluated within each ultrasound corpus.

6. RESULTS

6.1. Predicting the 9th frame

In Table 1, the performance of the ConvLSTM is compared with other predictors. As we used the same setting for the TJU data, some of the results obtained in [9] are also listed for comparison. The previous results include the MSE between the actual 9th frames and the output frame obtained from following predictors: the average of the preceding 8 video frames; the 8th frame alone; and the predicted 9th frame given by 3DCNN. The prediction of ConvLSTM neural network yields a lower MSE value than predictions given by all other methods.

As for the “Cross” condition, the task is essentially predicting only the tongue surface contour in the 9th frame. The

current results may suggest that 3DCNN could be more suitable for this specific subtask of predicting the snake-extracted contours of the 9th frame, and the less optimal performance of ConvLSTM might be due to the influence of other vocal tract features present in the ultrasound frames. However, though the ConvLSTM network does not perform as well as the 3DCNN, its predicted frames still yield far lower MSE than the baseline, which was obtained by computing the MSE between the 8th contour and the 9th contour. These results show that ConvLSTM also has good capacity to isolate the salient tongue contours in ultrasound frames.

Table 1. Mean MSE between the original frames and the predicted 9th frames by different predictors.

	WSJ0	TJU	Cross
Average	26.4	73.6 [9]	-
8 th image	22.6	40.0 [9]	279.5 [9]
3DCNN	21.3	32.6 [9]	154.9 [9]
ConvLSTM	13.2	28.1	169.2

The quality of the predicted frames is also assessed by the complex wavelet structural similarity index (CW-SSIM) [16], which has been shown to be particularly robust in measuring similarity between two ultrasound tongue images [17, 18]. In Table 2, the CW-SSIM performance of the convolutional LSTM neural network 9th image predictor is compared to the 3DCNN predictors. The CW-SSIM results again confirm that the predicted frames given by ConvLSTM are more similar to the original frames than those given by 3DCNN. Fig. 4 displays sample frames and the corresponding predicted frames from ConvLSTM outputs for visual comparison. Though predicted frames tend to be slightly more blurry than the original frames, these predicted images still contain highly similar tongue shapes and other vocal tract details to those original ultrasound images.

Table 2. Mean CW-SSIM between the original frames and the predicted 9th frames by different predictors.

	WSJ0	TJU
3DCNN	0.907	0.912
ConvLSTM	0.943	0.952

The training and test errors over time for ConvLSTM are displayed in Fig. 5. As shown in Fig. 5, the ConvLSTM model has high converging speed, high stability and good performance.

6.2. Predicting distant tongue motions

Table 3 and 4 display the outcomes of distant tongue motion prediction in consecutive ultrasound frames. The results show that both 3DCNN and ConvLSTM performed significantly better than the baseline prediction given by simple averaging




Fig. 4. Ultrasound frames and predicted frames in TJU corpus. The left two figures are the 9th frame and the predicted 9th frame respectively. The right two figures are the actual tongue contour and the predicted tongue contour for the same frame.




Fig. 5. ConvLSTM training and test losses over time for the WSJ0 (left) and TJU Data (right).

in either WSJ0 or TJU corpus, as measured by MSE and CW-SSIM. ConvLSTM consistently yields superior predictions in comparison with 3DCNN under all experimental conditions, further demonstrating its advantage in this particular task. As shown in Fig. 6, the predicted frames closely match the actual frames not only in tongue shapes but also in fine details.

It is also noted that predicting future tongue movements becomes increasingly challenging as the future tongue motion gradually diverges from the current motion. Consequently the prediction quality slightly decreases as the predicted frame gradually moves away from the input frames. The predictive performance on WSJ0 corpus is consistently better than that on TJU corpus, because the WSJ0 corpus was sampled at higher frame rates and tongue motion in a single frame tends to be more autocorrelated with its neighboring frames.

Table 3. MSE and CW-SSIM on TJU datasets for different predictors. The indices 10th or 11th indicate which frame is being predicted.

Predictor	MSE	CW-SSIM
Average-10 th	100.77	0.854
3DCNN-10 th	48.67	0.900
ConvLSTM-10 th	44.35	0.928
Average-11 th	116.09	0.841
3DCNN-11 th	66.35	0.881
ConvLSTM-11 th	61.38	0.904

Table 4. MSE and CW-SSIM on WSJ0 datasets for different predictors. The indices 10th or 11th indicate which frame is being predicted.

Predictor	MSE	CW-SSIM
Average-10 th	32.14	0.874
3DCNN-10 th	25.09	0.897
ConvLSTM-10 th	17.13	0.932
Average-11 th	35.28	0.869
3DCNN-11 th	28.83	0.873
ConvLSTM-11 th	20.87	0.913




Fig. 6. Ultrasound frames and predicted frames in TJU corpus. The left two figures are the 10th and 11th frames respectively. The right two figures are the corresponding predictions by ConvLSTM networks.

7. CONCLUSIONS

In this study, we demonstrate that ConvLSTMs are capable of predicting future frames in ultrasound videos at the pixel level with high accuracy. ConvLSTM consistently performs better than the 3DCNN and the baseline in predicting future frames, even when tongue motions are distant from the motions in input frames. Though the ConvLSTM shows slightly higher MSE in terms of predicting snake contours directly in comparison with 3DCNN, our method still demonstrates strong capacity to encode variations of tongue shapes in continuous ultrasound frames.

It is worth noting that the MSE loss function might be guiding the model to predict global vocal tract features in ultrasound frames rather than only the tongue shape, though neural networks do learn to predict the overall tongue shapes accurately. MSE also tends to bias model towards producing blurry images [9]. However, as the tongue shape constitutes the most critical speech-related feature in ultrasound frames, the next step would be to train models that only focus on the tongue shape features. Searching for alternative objective functions might further boost the model performance.

In the future, it would also be interesting to explore how prediction results vary across speakers and languages, and to what extent tongue motions are predictable. The current work may also have implications for practical applications, such as silent speech interface and articulatory motion synthesis.

8. REFERENCES

- [1] Jordan R Green and Yu-Tsai Wang, “Tongue-surface movement patterns during speech and swallowing,” *The Journal of the Acoustical Society of America*, vol. 113, no. 5, pp. 2820–2833, 2003.
- [2] Raymond D Kent and Kenneth L Moll, “Cinefluorographic analyses of selected lingual consonants,” *Journal of Speech, Language, and Hearing Research*, vol. 15, no. 3, pp. 453–473, 1972.
- [3] Hironori Takemoto, “Morphological analyses of the human tongue musculature for three-dimensional modeling,” *Journal of Speech, Language, and Hearing Research*, vol. 44, no. 1, pp. 95–107, 2001.
- [4] Maureen Stone, “A guide to analysing tongue motion from ultrasound images,” *Clinical linguistics & phonetics*, vol. 19, no. 6-7, pp. 455–501, 2005.
- [5] Min Li, Chandra Kambhamettu, and Maureen Stone, “Automatic contour tracking in ultrasound images,” *Clinical linguistics & phonetics*, vol. 19, no. 6-7, pp. 545–554, 2005.
- [6] Yan Ji, Licheng Liu, Hongcui Wang, Zhilei Liu, Zhibin Niu, and Bruce Denby, “Updating the silent speech challenge benchmark with deep learning,” *Speech Communication*, vol. 98, pp. 42 – 50, 2018.
- [7] Thomas Hueber, Guido Aversano, G Cholle, Bruce Denby, Gérard Dreyfus, Yacine Oussar, Pierre Roussel, and Maureen Stone, “Eigentongue feature extraction for an ultrasound-based silent speech interface,” in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP’07*. IEEE, 2007, vol. 1, pp. I–1245.
- [8] Jun Cai, Bruce Denby, Pierre Roussel, Gérard Dreyfus, and Lise Crevier-Buchman, “Recognition and real time performances of a lightweight ultrasound based silent speech interface employing a language model,” in *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [9] Chengrui Wu, Shicheng Chen, Guorui Sheng, Pierre Roussel, and Bruce Denby, “Predicting tongue motion in unlabeled ultrasound video using 3d convolutional neural networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5764–5768.
- [10] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu, “3d convolutional neural networks for human action recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, 2013.
- [11] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee, “Decomposing motion and content for natural video sequence prediction,” *arXiv preprint arXiv:1706.08033*, 2017.
- [12] Liang Zhang, Guangming Zhu, Peiyi Shen, Juan Song, Syed Afaq Shah, and Mohammed Bennamoun, “Learning spatiotemporal features using 3dcnn and convolutional lstm for gesture recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3120–3128.
- [13] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo, “Convolutional lstm network: A machine learning approach for precipitation nowcasting,” in *Advances in neural information processing systems*, 2015, pp. 802–810.
- [14] Swathikiran Sudhakaran and Oswald Lanz, “Learning to detect violent videos using convolutional long short-term memory,” in *Advanced Video and Signal Based Surveillance (AVSS), 2017 14th IEEE International Conference on*. IEEE, 2017, pp. 1–6.
- [15] Michael Kass, Andrew Witkin, and Demetri Terzopoulos, “Snakes: Active contour models,” *International journal of computer vision*, vol. 1, no. 4, pp. 321–331, 1988.
- [16] Mehul P Sampat, Zhou Wang, Shalini Gupta, Alan Conrad Bovik, and Mia K Markey, “Complex wavelet structural similarity: A new image similarity index,” *IEEE transactions on image processing*, vol. 18, no. 11, pp. 2385–2401, 2009.
- [17] Kele Xu, Tamás Gábor Csapó, Pierre Roussel, and Bruce Denby, “A comparative study on the contour tracking algorithms in ultrasound tongue images with automatic re-initialization,” *The Journal of the Acoustical Society of America*, vol. 139, no. 5, pp. EL154–EL160, 2016.
- [18] Kele Xu, Yin Yang, Maureen Stone, Aurore Jaumard-Hakoun, Clémence Leboullenger, Gérard Dreyfus, Pierre Roussel, and Bruce Denby, “Robust contour tracking in ultrasound tongue image sequences,” *Clinical linguistics & phonetics*, vol. 30, no. 3-5, pp. 313–327, 2016.