

Prediction Of Head Motion From Speech Waveforms With A Canonical-Correlation-Constrained Autoencoder

JinHong Lu, Hiroshi Shimodaira

Centre for Speech Technology Research, School of Informatics,
University of Edinburgh, United Kingdom

jinhong.l@sms.ed.ac.uk, H.Shimodaira@ed.ac.uk

Abstract

This study investigates the direct use of speech waveforms to predict head motion for speech-driven head-motion synthesis, whereas the use of spectral features such as MFCC as basic input features together with additional features such as energy and F0 is common in the literature. We show that, rather than combining different features that originate from waveforms, it is more effective to use waveforms directly predicting corresponding head motion. The challenge with the waveform-based approach is that waveforms contain a large amount of information irrelevant to predict head motion, which hinders the training of neural networks. To overcome the problem, we propose a canonical-correlation-constrained autoencoder (CC-CAE), where hidden layers are trained to not only minimise the error but also maximise the canonical correlation with head motion. Compared with an MFCC-based system, the proposed system shows comparable performance in objective evaluation, and better performance in subject evaluation.

Index Terms: head motion synthesis, speech-driven animation, deep canonically correlated autoencoder

1. Introduction

Head motion such as nodding and shaking is an important non-verbal communication channel in human-human communication. In addition to the head motion as nonverbal signals, Hadar *et al.* [1] have shown another type of head motion that is directly related to speech production. It is essential for animated talking heads to realise both types of natural head motion as well as lip-sync to make the avatar more human-like. Compared with lip-sync, the synthesis of head motion from audio speech is more challenging, since the link between speech and head motion is less clear, and not only speech, but also various factors such as emotion, intention, and stance are involved.

The present study considers a link between head motion and acoustic speech features, whose original representation is given as acoustic waveform signals, and seeks compact and efficient representation of speech features to predict corresponding head motion. Kuratate *et al.* [2] found that fundamental frequency (F0) and head motion had a correlation of 0.83 at sentence-level. Busso *et al.* [3] also confirmed a strong sentence-level correlation ($r=0.8$) between head motion and mel-frequency cepstral coefficients (MFCCs), where data was recorded for an actor reading the scripts of short sentences.

As we show in experiments, it is a different scenario in natural conversations, where there is a much larger degree of variation in head motion and we cannot find such strong correlations. A similar observation is reported for a dialogue corpus by Sadoughi *et al.* [4], where they have found a global canonical correlation analysis (global CCA) of 0.1931 between the original head movements and speech (F0 and energy).

In order to tackle the problem of a weak link between speech and head motion, other features and their combination have been explored. Ben-Youssef *et al.* [5] found that the articulatory features or EMA features that were estimated from speech were more useful to predict head motion. Ding *et al.* [6] examined LPC, MFCC, and filter bank (FBank) features and showed that FBank-based system outperformed MFCC-based one. Haag *et al.* [7] combined MFCC and EMA features to build bottleneck features, which were then fed to DNN-BLSTM to predict head motion. Greenwood *et al.* [8] proposed CVAE-BLSTM and used the decoder as a generative model to predict head motion, where the FBank features were used as the condition.

The purpose of using a combination of different features in the previous studies was to use richer information (e.g., prosodic features) to train models and predict head motion. Since all the acoustic features described above are derived from raw speech waveforms, it is natural to consider the original waveforms as the input to neural networks, so that we will be able to fully make use of the information in the original observations. So far, no one has investigated the use of original raw waveforms to predict head motion. This is mainly because of (1) the high dimensionality of raw waveform signals, which slows down the training of neural networks and requires high capacity in the hardware support; (2) a large amount of irrelevant information to predict head motion, which hinders the training of neural networks.

To overcome the problems of high dimensionality and irrelevant information, we propose a canonical-correlation-constrained autoencoder (CCCAE) to extract low-dimensional features from raw waveforms, where hidden layers are trained not only to minimise the error of encoding and decoding, but also maximise the canonical correlation with head motion. The extracted features of a low dimension are then fed to another neural network for regression to predict head motion. We show that the features obtained with the proposed approach are more useful for head-motion prediction than those with a standard autoencoder. We evaluate the new approach through comparisons with other acoustic features in terms of objective and subjective measures.

1.1. Relation to prior work

Using raw waveforms for acoustic modelling with neural networks is one of the active areas in automatic speech recognition [9, 10, 11, 12]. However, to the best of our knowledge, no one has investigated the use of raw waveforms for speech-driven head-motion synthesis, in which a set of two data streams, speech and head motion, is dealt with rather than a single stream of speech. Chandar *et al.* [13] and Wang *et al.* [14] have proposed the framework of correlational neural networks and deep

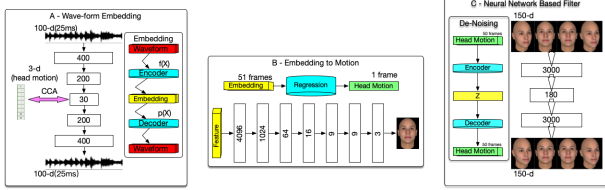


Figure 1: Overview of the proposed system comprised of three modules: (A) waveform embedding with CCCAE, (B) DNN-based head motion regression from the embedded features, (C) post filter with an autoencoder.

canonically correlated autoencoder (DCCA), respectively, to effectively model two data streams, and they applied the models to cross-language tasks and multi-view feature learning, where you can expect reasonably high correlations between two data streams. The present study is different in that the correlation between speech and head-motion features are much weaker, and our proposed model employs only one autoencoder whereas they employ two.

2. Proposed System

Our proposed system can be separated into three modules; (1) a canonical-correlation-constrained autoencoder (CCC AE) for compressing the high-dimensional waveform input to distributed embedding of low dimensions; (2) a regression model for predicting the head motion from the compressed embedding; (3) a post-filtering autoencoder for reconstructing smooth head motion. The overall framework of our proposed model is shown in Figure 1.

2.1. Waveform Embedding

The framework of autoencoder for a set of two data streams is proposed by Chandar *et al.* [13] and Wang *et al.* [14]. DCCA [14] consists of two autoencoders and optimises the combination of canonical correlation between the learned "bottleneck" representations and the reconstruction errors of the autoencoders. Since head motion is parameterised with a time series of rotation vectors of three dimensions in the present study, we do not need to use an autoencoder to reduce the dimensionality further. We thus employ a single autoencoder, in which hidden layers are trained in such a way as to not only minimise the reconstruction error, but also maximise the canonical correlation with head motion. Thus, instead of projecting the two features to a common subspace, we project raw waveforms to a subspace so that the embedded features are well correlated with head motion. We do not consider more advanced architectures such as VAE/CVAE [8, 15, 16], because standard AE is more effective in this task as the generative models, VAE/CVAE, is harder in the training.

We train the proposed CCCAE with the following objective function,

$$\text{Obj}_{\text{CCC AE}} = \sum_t \|\mathbf{X}_t - p(f(\mathbf{X}_t))\|^2 - \alpha \text{CCA}(f(\mathbf{X}), \mathbf{Y}) \quad (1)$$

where \mathbf{X}_t represents the input raw waveform vector at a time instance t to the encoder, $f(\cdot)$ represents the projection with the encoder, $p(\cdot)$ represents the reconstruction with the decoder, \mathbf{X} and \mathbf{Y} denote the whole sequences of waveform vectors and head motion vectors, respectively, and $\text{CCA}(\cdot)$ is the canonical correlation function. $\alpha \geq 0$ is the weighting factor, where

$\alpha = 0$ corresponds to a standard autoencoder with an MSE loss function.

2.2. Head motion regression

A simple feed-forward deep neural network is applied here for the regression from the waveform embedded features to head motion. We do not consider more complex models such as CNN and LSTM, because the present study focuses on a compact and efficient representation of speech features rather than the regression of head motion; and previous studies [6, 7] showed no large differences among the models. Accordingly, we also do not consider auto-regressive models such as WaveNet [17].

As is shown in Figure 1(B), a context window of ± 25 frames, which is equivalent to 525ms effective speech content, is employed to predict head motion parameters.

2.3. Post-filter

Since the output trajectories of the neural networks are noisy or discontinuous due to the nature of speech, it is common to apply a post-filter as post-processing to obtain smooth head motion trajectories for animation [7, 18, 19].

The common training procedure of the de-noising model, which comprises applying dropout/Gaussian noise to the clean data for recreating noisy data [20, 21], does not work with our model as the Gaussian noise method does not give the expected sudden and quick movements as they would naturally occur. The dropout method, on the other hand, drops one of the three trajectories of the head motion, and this strictly limits the movement, causing unnatural behaviour. Therefore, instead of removing the noise from the jerky head motion, we expect the de-noising filter to learn and know how the smooth head motion over a period should be. We assume a complete head motion in every consecutive 500ms [22] time frame, as the input, M_f , to the de-noising filter and the output are of the same length. We built a neural network based de-noising autoencoder following the architecture, trained with the "clean" data [23].

3. Experiments and Results

3.1. Dataset

We used the University of Edinburgh Speaker Personality and Mocap Dataset [24]. This dataset contains expressive dialogues of 13 native English semi-professional actors in extroverted and introverted speaking styles and the dialogues are non-scripted and spontaneous. For the purpose of our experiments, we selected data from one male (Subject A) and one female (Subject B). Six recordings (around 30 minutes) of each subject were used for training, two (around 10 minutes) for validation, and the remaining two (around 10 minutes) for evaluation, ignoring the differences in terms of the speaking style. Moreover, we have chosen the benchmark dataset, IEMOCAP [3], to validate our model. We selected the recordings of the first female subject from the dataset as Sadoughiet *al.* [19]. We trained our models for each subject. Note that speaker-dependent training is a common practice in speech-driven head motion synthesis [25, 6, 19].

3.1.1. Speech Features

Audio in the database was recorded with a headset microphone at 44.1 kHz with 32-bit depth and a MOTU 8pre mixer [26]. Separate recording channels were used for the two speakers and a synchronisation signal was recorded on a third channel

in the mixer. For the purpose of this work, the audio signal was downsampled to 4 kHz prior to feature extraction. Raw waveform vectors were extracted using 25 ms windows with 10 ms shifting, which resulted in 100 dimensions. 13 MFCCs feature is formed by combining 1 energy coefficient and 12 Mel-cepstral coefficients, using SPTK [27]. We also added their first and second-order derivatives, resulting in 39-MFCCs. Voicing probability and energy were computed using openS-MILE [28], and smoothed with a moving average filter with a window length of 10 frames. All the features were normalised in terms of variance for each dimension.

3.1.2. Head Motion Features

Movements of the head as a 3D rigid-body were recorded with the NaturalPoint Optitrack [29] motion capture system at a 100 Hz sampling rate. From the marker coordinates, rotation matrices for the head motion were computed using singular value decomposition [30], which were further converted to rotation vectors of three dimensions.

3.2. Experimental Setups

We conducted preliminary experiments to decide the depth and width of the models, which are shown in Figure 1. We tested different numbers of nodes, 15, 30, and 60, for the embedding layer of CCCAE, and decided to use 30 nodes based on the performance of the autoencoder. In training and objective evaluation, we only used the frames where the target speaker for head-motion prediction was speaking, so that the models learnt the relationship between speech and head motion properly. In subjective evaluation, we made use of all the input audio sequences to generate head motion parameters. The following notations are used in the rest experiments.

- $W_{v_{AE}}$: Embedded features extracted from the standard autoencoder (i.e., the output of proposed CCCAE with $\alpha = 0$)
- $W_{v_{CCCAE}}$: Embedded features extracted from the proposed CCCAE with $\alpha = 1$
- M_{MFCC} : Regression model trained with MFCC feature
- M_{AE} : Regression model trained with $W_{v_{AE}}$
- M_{CCCAE} : Regression model trained with $W_{v_{CCCAE}}$

M_{MFCC} , M_{AE} , and M_{CCCAE} use the same architecture in Figure 1(B) to predict head motion, while each model takes different feature vectors as input.

Training was conducted on a GPU machine and a multi-CPU machine with Tensorflow version 1.12 by mini-batch training using Adam optimisation (learning rate 0.0002) [31]. We also employed layer-wide pre-training [32].

3.3. Objective Evaluation

To measure the similarity between two sequences of vectors, we employed normalised mean-squared error (NMSE), where MSE is normalised by the variance of ground truth, and local CCA [7]. As opposed to global CCA, which calculates canonical correlations over the whole sequence, local CCA calculates CCA scores for every sub-sequence obtained with a time window and takes the average of the resulting scores. We used local CCA rather than global CCA, because head motion trajectories are not stationary and linear correlations rarely hold over long periods. We used a time window of 300 frames or 3 seconds.

In addition to the speech features described before, for comparison purposes, we also used F0+Energy (6 dimensions with

Table 1: *Local CCA between speech features and original head motion.*

Feature	Subject	CCA		
		Training	Valid	Test
F0+Energy	A	–	–	0.107
	B	–	–	0.117
	IEMOCAP	–	–	0.301
FBank	A	–	–	0.143
	B	–	–	0.157
	IEMOCAP	–	–	0.392
MFCC	A	–	–	0.238
	B	–	–	0.257
	IEMOCAP	–	–	0.784
waveform	A	–	–	0.186
	B	–	–	0.157
	IEMOCAP	–	–	0.658
$W_{v_{AE}}$	A	0.221	0.196	0.196
	B	0.110	0.135	0.176
	IEMOCAP	0.626	0.626	0.689
$W_{v_{CCCAE}}$	A	0.264	0.234	0.248
	B	0.220	0.240	0.266
	IEMOCAP	0.774	0.784	0.838

Table 2: *Comparison of different systems in terms of performance of head motion prediction, where NMSE and local CCA are calculated between predicted head motion and ground truth.*

System	Subject	Training		Test	
		NMSE	CCA	NMSE	CCA
M_{MFCC}	A	0.82	0.47	1.58	0.33
	B	0.55	0.57	1.71	0.34
	IEMOCAP	0.46	0.62	1.30	0.43
M_{AE}	A	1.01	0.18	1.07	0.20
	B	1.21	0.09	1.16	0.08
	IEMOCAP	0.98	0.01	1.11	0.00
M_{CCCAE}	A	0.58	0.42	1.52	0.29
	B	0.69	0.39	1.33	0.24
	IEMOCAP	0.68	0.61	1.14	0.37

delta and delta delta features) [4], FBank (27 dimensions of 26 filter-bank channels and log energy), and waveform (100 dimensions), which is the input to the proposed CCCAE.

3.3.1. CCA between speech features and original head motion

Before evaluating the performance of head-motion prediction from speech, we carried out a basic correlation analysis between speech features and head motion in terms of local CCA. Table 1 shows local CCA for each speech feature and for each subject. Note that CCA scores on training and validation sets are not shown for those features in which training is not involved. It can be found that F0+Energy gives the smallest, and MFCC and $W_{v_{CCCAE}}$ achieve the largest CCA scores with head motion. Compared to waveform, we can see a large improvement on the test set (by 33% for Subject A and 69% for Subject B) with $W_{v_{CCCAE}}$, whereas there is a small improvement with $W_{v_{AE}}$. Looking at the results of the benchmark dataset, our proposed feature has a improvement of 15% and 6.9% to waveform and MFCC respectively.

3.3.2. Evaluation of predicted head motion from speech

Based on the result of the basic analysis, we chose the three highest correlation features, MFCC, $W_{v_{AE}}$, and $W_{v_{CCCAE}}$ for

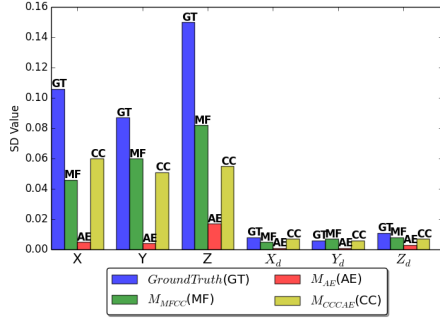


Figure 2: Comparison of predicted head motion trajectories with different models in terms of standard deviation of each rotation parameter (X , Y , Z) and its derivative, where values are averaged over the two subjects.

the evaluation of head-motion prediction. Table 2 shows the comparison of different systems, where the quality of predicted head motion was evaluated in terms of NMSE and local CCA with the ground truth (original head motion). We also computed local CCA between the ground truth and randomised sequences of another subject different from Subjects A and B to estimate a chance score for the two original head-motion sequences that are totally different and unsynchronised from each other and supposed to have no correlations. The estimated chance score for Subject A, Subject B and IEMOCAP respectively is $\rho_A = 0.16$, $\rho_B = 0.11$ and $\rho_{IEMOCAP} = 0.18$.

Although M_{AE} shows the lowest NMSE on the test set, it is just because the predicted head motion had little movement, which resulted in NMSE being close to 1.0. This is also reflected in the local CCA that, M_{AE} is worse than the chance score for both subjects. M_{CCCAE} gets performance comparable to M_{MFCC} in terms of NMSE. M_{MFCC} gets the highest local CCA. Overall, the quality of M_{MFCC} and M_{CCCAE} in the test dataset is higher than the chance score.

CCA captures only one aspect of similarity, i.e., linear correlations between two data streams, and it does not tell us how similar the two streams are in terms of other aspects such as dynamic range and smoothness, which we believe are also crucial factors in human perception. We thus calculated the standard deviation (SD) of each head motion trajectory and its derivative, i.e., velocity, whose result is shown in Figure 2. M_{AE} has the smallest SD in all trajectories and velocities, which confirms that M_{AE} has very little movement as mentioned above. The ground truth has the largest SD over trajectories. M_{MFCC} and M_{CCCAE} show comparable performance, but yet not close enough to the ground truth. Compared in terms of velocity, the ground truth, M_{MFCC} , and M_{CCCAE} are similar to each other. This indicates that they are likely to have the same level of smoothness.

3.4. Subjective Evaluation

We conducted a perceptual test using the Multiple Stimuli with Hidden Reference and Anchor (MUSHRA)[33]. We developed five test groups from Subject A, where each test group consisted of 3 randomly selected audio samples in the test set, and animations were created from each sample using 5 models: Ground Truth, M_{AE} , M_{CCCAE} , M_{MFCC} , and Anchor. Each animation lasts 8 – 12 seconds long (including speaking and listening frames). The anchor is created by selecting the original head

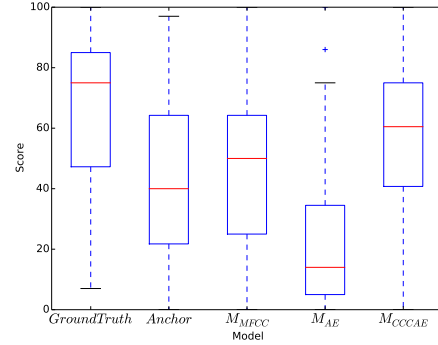


Figure 3: The Boxplot of the MUSHRA score for the Subject A's animation of each model - horizontal line indicates the median.

motion of another speaker with different utterances. This ensures that the anchor head motion has a natural behaviour, but it does not synchronise with the audio. The evaluation is performed such that every participant is assigned one test group and the animations of each test group are shown in a random order. Then, each participant watches each head-motion animation and gives a score, between 0 – 100, for each animation. A group of 20 participants were involved in this evaluation and they were asked to give a score to each animation according to the naturalness of the synthesised head motion.

The result is shown in Figure 3. M_{AE} scored the lowest among all including the anchor. We think the reason could be that as the predicted head motion with M_{AE} conducted a relatively minor movement, which may seem contrary to regular human beings' behaviour, from the participants' perspective. The anchor scored the second lowest as expected, participants were able to figure out the non-synchronicity between the head motion and audio. Compared between M_{MFCC} and M_{CCCAE} models, participants scored higher for M_{CCCAE} .

One explanation for M_{MFCC} achieving better in the objective, but lower in the subjective than M_{CCCAE} may be that the participants were affected by the active head motion generated by M_{MFCC} while the agent is listening, and they felt against natural human being. This also shows that objective is in things quantifiable, whereas subjective is one open to greater interpretation based on personal feeling[34].

4. Conclusions

In this paper, we have proposed an approach to create a highly correlated feature with head motion from raw waveform data using CCCAE. From the objective evaluations, we can conclude that (1) CCCAE enables creation of a more correlated feature ($W_{avCCCAE}$) with the head motion than W_{avAE} and other popular spectral features such as MFCC and FBank. (2) the M_{CCCAE} achieved the lowest NMSE in test dataset, although the local CCA is not the highest. (3) the analysis based on SD shows that M_{MFCC} and M_{CCCAE} are comparable performance. (2) and (3) indicate that $W_{avCCCAE}$ is capable of being used in achieving state-of-the-art results for predicting natural head motion with the advantage of the CCCAE. (4) MUSHRA test shows that excluding the ground truth, participants preferred to choose the animation generated by M_{CCCAE} over the others. Overall, our M_{CCCAE} shows better performance than M_{MFCC} .

5. References

- [1] U. Hadar, T. Steiner, E. Grant, and F. Rose, "Head movement correlates of juncture and stress at sentence level," *Language and Speech*, vol. 26, no. 2, pp. 117–129, 1983.
- [2] T. Kuratate, K. G. Munhall, P. E. Rubin, E. Vatikiotis-Bateson, and H. Yehia, "Audio-visual synthesis of talking faces from speech production correlates," in *Eurospeech'99*, vol. 3, 1999, pp. 1279–1282.
- [3] C. Busso and S. Narayanan, "Interrelation Between Speech and Facial Gestures in Emotional Utterances: A Single Subject Study," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2331–2347, November 2007.
- [4] N. Sadoughi and C. Busso, "Speech-driven Animation with Meaningful Behaviors," *CoRR*, vol. abs/1708.01640, 2017. [Online]. Available: <http://arxiv.org/abs/1708.01640>
- [5] A. Ben Youssef, H. Shimodaira, and D. Braude, "Speech driven talking head from estimated articulatory features," in *Proc. ICASSP*, 2014, pp. 4573–4577.
- [6] C. Ding, L. Xie, and P. Zhu, "Head motion synthesis from speech using deep neural networks," *Multimedia Tools and Applications*, vol. 74, no. 22, pp. 9871–9888, November 2015. [Online]. Available: <https://doi.org/10.1007/s11042-014-2156-2>
- [7] K. Haag and H. Shimodaira, "Bidirectional LSTM Networks Employing Stacked Bottleneck Features for Expressive Speech-Driven Head Motion Synthesis," in *Intelligent Virtual Agents*, 2016, pp. 198–207.
- [8] D. Greenwood, S. Laycock, and I. Matthews, "Predicting head pose from speech with a conditional variational autoencoder," in *Interspeech*, 08 2017, pp. 3991–3995.
- [9] T. N. Sainath, R. J. Weiss, A. W. Senior, K. W. Wilson, and O. Vinyals, "Learning the speech front-end with raw waveform CLDNNs," in *INTERSPEECH*, 2015.
- [10] P. Ghahremani, V. Manohar, D. Povey, and S. Khudanpur, "Acoustic Modelling from the Signal Domain Using CNNs," in *INTERSPEECH*, 2016, pp. 3434–3438.
- [11] Z. Tüske, R. Schlüter, and H. Ney, "Acoustic Modeling of Speech Waveform Based on Multi-Resolution, Neural Network Signal Processing," in *ICASSP*, 2018, pp. 4859–4863.
- [12] P. von Platen, C. Zhang, and P. Woodland, "Multi-Span Acoustic Modelling using Raw Waveform Signals," *ArXiv*, vol. abs/1906.11047, 2019.
- [13] S. Chandar, M. M. Khapra, H. Larochelle, and B. Ravindran, "Correlational Neural Networks," *CoRR*, vol. abs/1504.07225, 2015. [Online]. Available: <http://arxiv.org/abs/1504.07225>
- [14] W. Wang, R. Arora, K. Livescu, and J. A. Bilmes, "On Deep Multi-View Representation Learning: Objectives and Optimization," *CoRR*, vol. abs/1602.01024, 2016. [Online]. Available: <http://arxiv.org/abs/1602.01024>
- [15] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *CoRR*, vol. abs/1312.6114, 2014.
- [16] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 3483–3491. [Online]. Available: <http://papers.nips.cc/paper/5775-learning-structured-output-representation-using-deep-conditional-generative-models.pdf>
- [17] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: A Generative Model for Raw Audio," *CoRR*, vol. abs/1609.03499, 2016.
- [18] C. Ding, P. Zhu, and L. Xie, "BLSTM neural networks for speech driven head motion synthesis," in *INTERSPEECH*, 2015.
- [19] N. Sadoughi and C. Busso, "Novel Realizations of Speech-Driven Head Movements with Generative Adversarial Networks," in *ICASSP*, April 2018, pp. 6169–6173.
- [20] P. Ghosh, J. Song, E. Aksan, and O. Hilliges, "Learning Human Motion Models for Long-term Predictions," *CoRR*, vol. abs/1704.02827, 2017. [Online]. Available: <http://arxiv.org/abs/1704.02827>
- [21] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion," *J. Mach. Learn. Res.*, vol. 11, pp. 3371–3408, December 2010. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1756006.1953039>
- [22] G. Hofer and H. Shimodaira, "Automatic head motion prediction from speech data," in *INTERSPEECH*, 2007.
- [23] J. Lu and H. Shimodaira, "A neural network based post-filter for speech-driven head motion synthesis," *arXiv e-prints*, p. arXiv:1907.10585, July 2019.
- [24] K. Haag and H. Shimodaira, "The University of Edinburgh Speaker Personality and MoCap Dataset," in *Facial Analysis and Animation Proceedings, Vienna*, 2015.
- [25] C. Busso, Z. Deng, U. Neumann, and S. Narayanan, "Natural head motion synthesis driven by acoustic prosodic features: Virtual Humans and Social agents," *Computer Animation and Virtual Worlds*, vol. 16, pp. 283–290, July 2005.
- [26] "MOTU," <http://www.motu.com/>.
- [27] "Speech Signal Processing Toolkit (SPTK)," <http://sptk.sourceforge.net>.
- [28] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: The Munich Versatile and Fast Open-source Audio Feature Extractor," in *Proceedings of the 18th ACM International Conference on Multimedia*, 2010, pp. 1459–1462. [Online]. Available: <http://doi.acm.org/10.1145/1873951.1874246>
- [29] "Naturalpoint Optitrack," <http://www.naturalpoint.com/optitrack>.
- [30] I. Soderkvist and P. Wedin, "Determining the movements of the skeleton using well-configured markers," *Journal of Biomechanics*, vol. 26, pp. 1473–1477, January 1994.
- [31] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [32] S. Takaki and J. Yamagishi, "A deep auto-encoder based low-dimensional feature extraction from FFT spectral envelopes for statistical parametric speech synthesis," in *ICASSP*, March 2016, pp. 5535–5539.
- [33] "Method for the subjective assessment of intermediate quality level of coding systems," recommendation ITU-R BS.1534 https://www.itu.int/dms_pubrec/itu-r/rec/bs/R-REC-BS.1534-1-200301-S!!PDF-E.pdf.
- [34] L. Leahu, S. Schwenk, and P. Sengers, "Subjective objectivity: Negotiating emotional meaning," in *DIS '08*, 01 2008, pp. 425–434.