論文 / 著書情報
Article / Book Information

Title	Detecting Alzheimer's Disease Using Gated Convolutional Neural Network from Audio Data
Authors	Tifani Warnita, Nakamasa Inoue, Koichi Shinoda
Citation	Proc. Interspeech 2018, , , pp. 1706-1710, 1713
Pub. date	2018, 9
Copyright	(c) 2018 International Speech Communication Association, ISCA
DOI	http://dx.doi.org/10.21437/Interspeech.2018-1713

Detecting Alzheimer's Disease Using Gated Convolutional Neural Network from Audio Data

Tifani Warnita, Nakamasa Inoue, Koichi Shinoda

Tokyo Institute of Technology, Tokyo, Japan

{tifani,inoue}@ks.cs.titech.ac.jp, shinoda@c.titech.ac.jp

Abstract

We propose an automatic detection method of Alzheimer's diseases using a gated convolutional neural network (GCNN) from speech data. This GCNN can be trained with a relatively small amount of data and can capture the temporal information in audio paralinguistic features. Since it does not utilize any linguistic features, it can be easily applied to any languages. We evaluated our method using Pitt Corpus. The proposed method achieved the accuracy of 73.6%, which is better than the conventional sequential minimal optimization (SMO) by 7.6 points.

Index Terms: Alzheimer's disease, computational paralinguistics, convolutional neural network, gating mechanism

1. Introduction

Alzheimer's Disease (AD) is the most common cause of dementia [1], a neurodegenerative disease strongly related to the reduced functionality or even death of neurons in the central nervous system [2]. As the result of aging society, we face an increasing number of people being affected by AD [3] which is estimated to double every 20 years [4]. The most noticeable symptom of AD patients is the memory loss [5] such as in recalling experiences, which results in poor narrative memory [6]. They also often become apathy and get depressed easily [1].

At present, there is no clear protocol on how to detect AD not only in an effective but also accurate way [1]. The most common approach is to monitor the patients, to carefully examine the medical history of the patients, to conduct some cognitive tests (i.e. a picture description task, a naming task), mental status, and mood test, as well as to take their brain images. The careful diagnosing process can take several days or even weeks which might be very cumbersome. Early prediction of AD actually can help its patients to preserve their cognitive functions [7]. Some of the causes are treatable and the patients are sometimes fully recovered [8]. Automatic detection of AD in its early phase has been strongly demanded.

To date, there have been a lot of studies for automatically detecting AD patients. Most of them used linguistic information [3, 9] since language impairment can be one of the most prominent cues of AD patients [3]. This leads to the difficulty in applying the approach to different languages, especially to low-resource languages. In this study, we propose a non-linguistic approach for detecting AD using acoustic features from speech data which is expected to give the flexibility in applying it in various languages. Inspired by numerous successes with deep learning for paralinguistic tasks such as for emotion recognition [10, 11], we employ convolutional neural networks with a gating mechanism.

2. Related Works

Linguistic information had been widely utilized in order to automatically detect AD patients [12, 13, 9]. For instance, Wankerl et al. (2017) [9] employed a pure linguistic approach based on n -gram on Pitt Corpus [14], in which subjects undergo a picture description task. From the study, they found that patients often not only uttered incomplete phrases but also interrupted others. These degraded the intelligibility of their speech [15].

Some studies utilized both linguistic features and acoustic features to detect people having AD [15, 3, 16]. Fraser et al. (2016) [3] utilized multilinear logistic regression and selected 35 top-ranked features out of 370 features by using Pitt Corpus [14]. Khodabakhsh et al. (2015) [15] used the recordings of conversational speech of 32 AD patients and 51 control subjects for this combination. Weiner et al. (2016) [17] employed only acoustic information from German conversational recordings. However, the study relied on the transcription for calculating some features such as silence-to-word ratio and word rate.

In contrast with these studies, we focused on non-linguistic approach where only speech audio features of the subjects are utilized. It can be easily applied to other languages and expected to be more robust against environmental noise and channel differences. As features, we used the paralinguistic features that frequently used in emotion recognition task since previous studies showed that AD people tend to show emotional prosodic impairment [18]. Numerous approaches for detecting emotions emerged from people have been carried out. Schuller et al. (2009) [19] provided the baseline for emotion recognition in the INTERSPEECH 2009 Emotion Challenge, which uses sequential minimal optimization (SMO). Huang et al. (2014) [10] employed deep learning based approach in which convolutional layers were employed. Keren and Schuller (2016) [11] used not only convolutional but also recurrent layers.

3. Database

In this study, we used the picture description task session data of Pitt Corpus [14], which is a part of DementiaBank, a multimedia database for studying people having dementia. In the picture description task, patients were asked to describe what happens in a picture, Cookie Theft Picture of the Boston Diagnostic Aphasia Examination [20]. Pitt Corpus consists of speech data and their transcription from 244 control (healthy) subjects as well as 309 patients having dementia such as mild cognitive impairment (MCI), vascular dementia, and AD. In this study, we used only the data from AD patients and from patients suspected of having AD (probable AD). It should be noted that, even though we used the same database as in [3, 9], the number of the subjects in our study was slightly different compared to theirs according to the following three reasons. First, the size of the database has increased over time. Second, we ex-

cluded speech data with overlapped sounds from the other interview sessions. Third, we only used data with both audio data and transcription information. Consequently, the number of sessions is 488 (255 AD, 233 control) from 267 subjects (169 AD, 98 control). Similar to [3], due to a limited amount of data, we employed a 10-fold cross-validation scheme for our evaluation. We designed the ten subsets so that no subject appeared in more than one subset in order to avoid any speaker dependencies.

In this study, we performed four preprocessing stages. First, we decrease the sampling rate into 16kHz. After that, we normalize each audio signal using the average value of decibels relative to full scale (dBFS) in the data. Then, we segmented the audio data of each subject into utterances according to the transcription information. We obtained the total of 6267 utterances (3276 AD, 2991 control). Finally, we added 10ms at the beginning and at the end of each utterance segment with 15ms fade-in and fade-out.

4. Features

In this study we used openSMILE [21] to extract acoustic features from Pitt Corpus. OpenSMILE consists of several different configurations of acoustic feature extraction. It has been mainly used for emotion recognition but we believe it is also effective for our purpose. While AD patients can easily get depressed, anxious, or even upset [1], they find it difficult to express their emotions in prosodies such as tempo alteration and powerful intonation [18]. Based on this finding, we used the following feature sets.

1. INTERSPEECH 2009 Emotion Challenge Features (IS09) [19]

In this feature set, there are 16 types of low-level descriptors (LLDs) extracted from the frame level. The LLDs include several prosodic features such as pitch and energy which are expected to represent the emotional pattern of the speech. The delta coefficients are also calculated hence producing the total of 32 LLDs. In order to get the utterance-level features from the LLDs, 12 functionals (e.g. the values of minimum and maximum, mean, and range) are applied to each LLDs. As a result, 384 features are extracted from one utterance.

2. INTERSPEECH 2010 Paralinguistic Challenge Features (IS10) [22]

The additional LLDs to IS09 are PCM loudness, eight log Mel frequency band (0-7), eight line spectral pairs (LSP) frequency (0-7), F0 envelope, voicing probability, jitter local, jitter DPP, and shimmer local. These features are expected to give broader coverage of paralinguistic events in speech. In addition, more MFCC features are extracted (0-14 compared to 1-12). Finally, we get 76 LLDs for one frame and the total of 1582 features for one utterance.

3. INTERSPEECH 2011 Speaker State Challenge Features (IS11) [23]

Compared to the previous feature sets, IS11 provides the derived loudness measure and the employment of Relative Spectral Analysis (RASTA)-style filtered auditory spectra resulting in 118 LLDs for the frame-level features. The total number of features in one utterance is 4368.

4. INTERSPEECH 2012 Speaker Trait Challenge Features (IS12) [24]

In this feature set, some of the LLDs added compare to IS11 include harmonic-to-noise ratio (HNR), spectral harmonicity, and psychoacoustic spectral sharpness resulting in 120 LLDs for the frame-level features. After being applied by functionals, the total number of features in one utterance is 6125 features.

5. Method

5.1. Convolutional Neural Network

In this study, assuming that temporal features are well represented in the frame-level features, we employed a Convolutional Neural Network (CNN) [25] as a classifier. A CNN was inspired by the structure of animal visual cortex for perceiving lights [26] and has yielded supreme results in an abundance of tasks in the past few years [27]. Furthermore, it needs a relatively small amount of training data compared to the other networks since it has a much smaller number of connection weights [28].

The convolution operation over the input aims to extract the temporal information by sliding through the use of a kernel (filter). In our study, we feed utterance segment features $X \in \mathbb{R}^{F \times T}$ into the CNN, where F and T represent the dimension of LLDs and the number of time frames, respectively. The size of the sliding window is $\mathbb{R}^{F \times N}$ where N is the window length of the kernel in the time axis. The convolution operation between the kernel and the input will produce a scalar output. Output from the convolution layer at time i , $i = n + 1, \dots, T$ is then defined as,

$$y_i = g \left(\sum_{f=1}^F \sum_{n=1}^N w_{f,n} x_{f,i-n} + b \right), \quad (1)$$

where b and $g(\cdot)$ denotes the bias and activation function respectively. $x_{c,d}$ is the (c, d) element of the input X , and $w_{c,d}$ is the (c, d) element of the weight matrix W . Both W and b are the learnable parameters of the network that we train. When the kernel is sliding through the input matrix over the time dimension, we multiply the overlapping element of the two matrices as in Eq. 1. We employed a rectified linear unit (ReLU) [29] as the activation function g . This network is also referred as Time-Delay Neural Network (TDNN) [30].

Since the input dimension for CNN should be fixed, we set the segment length as that of the longest utterance segment in the dataset. After that, we applied zero padding for the rest of the utterance segment. We use the number of kernels $K = 64$ and the window length of $N = 3$. Accordingly, the calculation of every patch of the input will produce the complete output $Y \in \mathbb{R}^{M \times K}$ where M is the number of time frames, $T - N + 1$.

We added batch normalization before each activation function [31]. We also used random weight initialization for each convolution layer. After each convolution layer, we put a max-pooling layer [32].

The output of the last convolution layer is flattened into one feature vector. For example, the flattened output $Y \in \mathbb{R}^{M \times N}$ will produce a vector $Z \in \mathbb{R}^O$ where $O = M \times N$. This vector became the input to a fully-connected layer with activation function ReLU consisting of 256 hidden neurons. We also employed batch normalization and initialized the weight matrix as random uniform. Furthermore, we also applied dropout with the value of 0.5 before the output layer. The output layer consists of one hidden neuron with a sigmoid function. We used binary cross-entropy as the loss function and Adam [33] as the optimizer. We trained the network with the pair of LLDs of each

Figure 1: GCNN with $depth = 1$. The kernel window in the convolutional layer is colored in gray.

utterance segment and its corresponding binary label.

5.2. Gating Mechanism

In addition to the standard deep convolutional neural network, we introduce a gating mechanism after each convolution. The resulting network is called Gated Convolution Neural Network (GCNN). A gate represents an information controller which manages the information that flows into the succeeding layer. The gate has the ability to prevent the gradient from being vanished during backpropagation [34]. Recently, it has been often used in convolutional neural networks with various objectives such as for conditional image modeling [35], language modeling [34], and speech synthesis [36]. A previous study [34] showed that gated linear unit (GLU) outperformed gated tanh unit (GTU) which is used in [35]. Therefore, we employed GLU in our study.

Similar as in Section 5.1, we give the input of extracted LLDs features of each utterance segment $X \in \mathbb{R}^{T \times F}$. We used a sigmoid function as the activation function g , which is multiplied by a linear gate. Eq. 1 is modified as:

$$y_i = \left(\sum_{f=1}^F \sum_{n=1}^N v_{f,n} x_{f,i-n} + e \right) \cdot g \left(\sum_{f=1}^F \sum_{n=1}^N w_{f,n} x_{f,i-n} + b \right), \quad (2)$$

Figure 2: Comparison of IS09, IS10, IS11, and IS12 feature sets using standard CNN with sb and ub denote the subject-level and utterance-level based classification respectively.

where $v_{c,d}$ represents the element of V at position (c, d) . The $V \in \mathbb{R}^{K \times F}$ and $e \in \mathbb{R}$ are the kernel weight matrix and bias for the linear gate. For GCNN, we used $N = 2$ and applied the same padding.

After that, we halved the output length in a max-pooling layer [32]. Figure 1 shows the visualization of our GCNN with the $depth = 1$ where the $depth$ represents the number of gated convolution layers. In the figure, one gated convolution layer lies between the dotted horizontal line which followed by the max-pooling layer. Deeper networks consist of more gated convolution layers. We also defined the same layers after the max-pooling layer as in Section 5.1.

5.3. Framework Overview

While we need to classify each subject based on his/her whole data, we performed the classification for each utterance instead. This is based on our assumption that the information about a patient having AD or not can still be obtained from a smaller but appropriate length of a segment. After the utterance-based classification is performed, we make the final verdict for each subject based on the proportion of each class; we classified a subject into AD if he/she has AD percentage above the half. The utterance-level subject classification is expected to give more detailed information about the symptoms while most of the previous studies conducted the subject-level classification [15, 3, 12, 13, 9, 16].

6. Experimental Results

We used the average accuracy from the 10-fold cross-validation scheme on Pitt Corpus [14] (see Section 3 for details) for evaluating our method. Our first experiment was the employment of the standard deep CNN without gates on the four feature sets, IS09, IS10, IS11, and IS12. Figure 2 shows the classification result of AD or non-AD with different numbers of hidden layers (1, 2, 3, 4, 6, 8, 10). The performance result given is from the subject-level classification after majority voting from the utterance-level classification. Since there is no established

Table 1: Comparison of the employment of the standard CNN and GCNN by using IS10 feature set represented in the average accuracy (%) from 10-fold cross-validation.

Depth	Standard CNN		Linear GCNN	
	Utterance	Subject	Utterance	Subject
1	64.2	66.0	62.2	66.2
2	64.2	66.6	62.6	68.7
3	64.2	69.2	61.9	66.4
4	64.9	68.7	63.3	68.9
6	65.5	68.6	65.1	72.2
8	66.1	69.0	66.3	73.6
10	65.2	70.4	65.2	69.8

Table 2: Confusion matrix depicting the prediction results from all folds using IS10 feature set and GCNN with the depth of 8.

		Actual			
		Utterance-Level		Subject-Level	
		AD	Control	AD	Control
Pred	AD	2340	1213	189	65
	Control	936	1778	66	168

baseline for this database and the number of instances used are different from one experiment to another experiment, we give the result of the four feature sets with baseline methods, which is the subject-level classification using sequential minimal optimization (SMO) [19, 22]. In Figure 2, they are marked by a star symbol (*).

From Figure 2, the best result is obtained when we use 10-layer CNN with IS10 feature set which is better than SMO by 2.4 points. Furthermore, we can see that the use of CNN with both IS11 and IS12 could not yield better result compared to using SMO while it can improve the performance of using IS09 and IS10. We can also see that the IS10 feature set outperformed the rest of the feature sets when we use CNN.

When we compare between IS09 and IS10, we can see that IS10 covers more features in the paralinguistic aspect of speech as it was used as the age-gender and level-of-interest classification. The performance of the feature sets IS09 was worse than that of IS10 especially when the subjects did not have any specific emotion (neutral). Some emotions appeared at the beginning before the subject begins to describe the image, in the middle when the subject begins to confuse with the things he/she wants to describe (laughs), and at the end after completing the task.

Noticeable differences between IS09 and IS10 include the use of voicing probability in which might more represent the sound-silence pattern in the subjects. Further, jitter and shimmer in IS10 might give more representation of the hesitation rate in the subjects. Those LLDs also appear in both IS11 and IS12. However, the higher dimension of the two feature sets might be too big for the input of the CNN.

Next, we investigated the effectiveness of the gating mechanism. The experiment was carried out by using IS10 feature set. The comparison of the standard CNN and the gating mechanism is shown in Table 2. From the table, we can see that the employment of linear gate improved the average accuracy from the 10-fold cross-validation scheme into 73.6%. We can see that the gating mechanism yields better results.

Furthermore, we also present the confusion matrix of the

Figure 3: GCNN with different utterance length.

best model showing the aggregated values from the ten folds in Table 1, both from the utterance-level and the subject-level. The confusion matrix shows the balance performance in predicting both of AD and Control subject. We can also see that predicting the individual utterance of the subject can be considered as a difficult task. However, combining several utterances for one subject can improve the subject-level prediction result.

Lastly, we investigated the importance of utterance length information in the classification performance. We tried a set of different segment length L , which are 500ms, 1000ms, 2000ms, 4000ms, and also 4295ms which is the segment length of the longest utterance in the dataset. In this case, we segmented each subject data into segments with a predetermined length L without taking into consideration the oracle utterance length. Accordingly, the zero padding is added only for the last utterance segment of the subject if its length is less than L . The experiment was carried out by using the best scheme from the previous experiment, which is GCNN. We only tried to use the gated CNN with the number of layers as 6, 8, and 10.

As depicted in Figure 3, we can see that shorter segment length yields worse results. However, we can still get close results using the utterance length of 4000ms (69.1%, 70.8%, and 69.8% for 6, 8, and 10 layers respectively) compared to the cases when we use the oracle utterance length (72.2%, 73.6%, and 69.8% for 6, 8, and 10 layers respectively). For 10 layers, we got the similar result between using the utterance length of 4000ms and the oracle utterance length with 69.8%. This result suggests that we can use the approach even if we do not have the transcription.

7. Conclusion

We present our study in the non-linguistic approach for detecting AD by utilizing only the speech audio data. The employment of GCNN yielded the best result of 73.6%. Since it does not utilize linguistic information, we can easily apply it to low-resource languages.

There are still a lot of remaining tasks. In the near future, we will evaluate our current approach on different languages especially on low-resource languages. The study on how to early detect AD patients is also a major concern. Other possible directions include estimating the severity of the disease and also evaluating its temporal change.

8. Acknowledgements

This work was supported by JSPS KAKENHI 16H02845 and by JST CREST Grant Number JPMJCR1687, Japan.

9. References

- [1] A. Association, "2017 Alzheimer's disease facts and figures," vol. 13, 2017.
- [2] T. Yacoubian, "Neurodegenerative disorders: Why do we need new therapies?" pp. 1–16, 2017.
- [3] K. C. Fraser, J. A. Meltzer, and F. Rudzicz, "Linguistic features identify Alzheimer's disease in narrative speech," *Journal of Alzheimer's Disease*, vol. 49, no. 2, pp. 407–422, 2000.
- [4] C. P. Ferri, M. Prince, C. Brayne, H. Brodaty, L. Fratiglioni, M. Ganguli, K. Hall, K. Hasegawa, H. Hendrie, Y. Huang *et al.*, "Global prevalence of dementia: a Delphi consensus study," *The lancet*, vol. 366, no. 9503, pp. 2112–2117, 2005.
- [5] A. Burns and S. Iliffe, "Alzheimer's disease," *BMJ*, vol. 338, 2009.
- [6] E. T. Prud'hommeaux and B. Roark, "Extraction of narrative recall patterns for neuropsychological assessment," in *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [7] H. Martono, *Buku Ajar Boedhi-Darmojo: Geriatri (Ilmu Kesehatan Usia Lanjut) [Boedhi-Darmojo Textbook: Geriatrics (Elderly People Health Science)]*. Balai Penerbit FKUI, 2000.
- [8] T. Manjari, V. Deepti *et al.*, "Reversible dementias," *Indian Journal of Psychiatry*, vol. 51, no. 5, pp. 52–55, 2009.
- [9] S. Wankerl, E. Nöth, and S. Evert, "An N -gram based approach to the automatic diagnosis of Alzheimer's disease from spoken language," in *Proc. INTERSPEECH 2017*, 2017.
- [10] Z. Huang, M. Dong, Q. Mao, and Y. Zhan, "Speech emotion recognition using CNN," in *ACM Multimedia*, 2014.
- [11] G. Keren and B. Schuller, "Convolutional RNN: an enhanced model for extracting features from sequential data," *2016 International Joint Conference on Neural Networks (IJCNN)*, pp. 3412–3419, 2016.
- [12] V. C. Zimmerer, M. Wibrow, and R. A. Varley, "Formulaic language in people with probable Alzheimer's disease: a frequency-based approach," *Journal of Alzheimer's Disease*, vol. 53, no. 3, pp. 1145–1160, 2016.
- [13] S. O. Orimaye, J. S. Wong, K. J. Golden, C. P. Wong, and I. N. Soyiri, "Predicting probable Alzheimer's disease using linguistic deficits and biomarkers," *BMC bioinformatics*, vol. 18, no. 1, p. 34, 2017.
- [14] J. T. Becker, F. Boiler, O. L. Lopez, J. Saxton, and K. L. McGonigle, "The natural history of Alzheimer's disease: description of study cohort and accuracy of diagnosis," *Archives of Neurology*, vol. 51, no. 6, pp. 585–594, 1994.
- [15] A. Khodabakhsh, F. Yesil, E. Guner, and C. Demiroglu, "Evaluation of linguistic and prosodic features for detection of Alzheimer's disease in Turkish conversational speech," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2015, no. 1, p. 9, 2015.
- [16] R. Sadeghian, J. D. Schaffer, and S. A. Zahorian, "Speech processing approach for diagnosing dementia in an early stage," *Proc. INTERSPEECH 2017*, pp. 2705–2709, 2017.
- [17] J. Weiner, C. Herff, and T. Schultz, "Speech-based detection of Alzheimer's disease in conversational german," pp. 1938–1942, 2016.
- [18] G. Tosto, M. Gasparini, G. Lenzi, and G. Bruno, "Prosodic impairment in Alzheimer's disease: assessment and clinical relevance," *The Journal of neuropsychiatry and clinical neurosciences*, vol. 23, no. 2, pp. E21–E23, 2011.
- [19] B. Schuller, S. Steidl, and A. Batliner, "The INTERSPEECH 2009 emotion challenge," *Proc. INTERSPEECH 2009*, pp. 312–315, 2009.
- [20] E. Kaplan, H. Goodglass, S. Weintraub, and S. Brand, *Boston Naming Test*. Lea & Febiger, 1983.
- [21] F. Eyben, F. Weninger, F. Gross, and B. Schuller, "Recent developments in opensmile, the munich open-source multimedia feature extractor," pp. 835–838, 2013.
- [22] B. Schuller, S. Steidl, A. Batliner, F. Burkhardt, L. Devillers, C. Müller, and S. Narayanan, "The INTERSPEECH 2010 paralinguistic challenge," *Proc. INTERSPEECH 2010*, pp. 2794–2797, 2010.
- [23] B. Schuller, S. Steidl, A. Batliner, F. Schiel, and J. Krajewski, "The INTERSPEECH 2011 speaker state challenge," *Proc. INTERSPEECH 2011*, 2011.
- [24] B. Schuller, S. Steidl, A. Batliner, E. Nöth, A. Vinciarelli, F. Burkhardt, R. v. Son, F. Weninger, F. Eyben, T. Bocklet *et al.*, "The INTERSPEECH 2012 speaker trait challenge," *Proc. INTERSPEECH 2012*, 2012.
- [25] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural computation*, vol. 1, no. 4, pp. 541–551, 1989.
- [26] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai *et al.*, "Recent advances in convolutional neural networks," *Pattern Recognition*, vol. 77, pp. 354–377, 2017.
- [27] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [28] Y. Bengio *et al.*, "Learning deep architectures for ai," *Foundations and trends in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.
- [29] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature*, vol. 405, no. 6789, p. 947, 2000.
- [30] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," in *Readings in speech recognition*. Elsevier, 1990, pp. 393–404.
- [31] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," pp. 448–456, 2015.
- [32] Y. Zhou and R. Chellappa, "Computation of optical flow using a neural network," vol. 27, pp. 71–78, 1988.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [34] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," pp. 933–941, 2017.
- [35] A. v. d. Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, "Conditional image generation with pixelcnn decoders," pp. 4790–4798, 2016.
- [36] A. v. d. 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," *arXiv preprint arXiv:1609.03499*, 2016.