Language learning using Speech to Image retrieval

Danny Merx, Stefan L. Frank, Mirjam Ernestus

Centre for Language Studies, Radboud University, Nijmegen, The Netherlands

d.merkx@let.ru.nl, s.frank@let.ru.nl, m.ernestus@let.ru.nl

Abstract

Humans learn language by interaction with their environment and listening to other humans. It should also be possible for computational models to learn language directly from speech but so far most approaches require text. We improve on existing neural network approaches to create visually grounded embeddings for spoken utterances. Using a combination of a multi-layer GRU, importance sampling, cyclic learning rates, ensembling and vectorial self-attention our results show a remarkable increase in image-caption retrieval performance over previous work. Furthermore, we investigate which layers in the model learn to recognise words in the input. We find that deeper network layers are better at encoding word presence, although the final layer has slightly lower performance. This shows that our visually grounded sentence encoder learns to recognise words from the input even though it is not explicitly trained for word recognition.

Index Terms: speech recognition, multimodal embeddings, computational linguistics, deep learning

1. Introduction

Most computational models of natural language processing (NLP) are based on written language; machine translation, sentence meaning representation and language modelling to name a few (e.g. [1, 2]). Even if the task inherently involves speech, such as in automatic speech recognition, models require large amounts of transcribed speech [3]. Yet, humans are capable of learning language from raw sensory input, and furthermore children learn to communicate long before they are able to read. In fact, many languages have no orthography at all and there are also languages of which the writing system is not widely used by its speakers. Text-based models cannot be used for these languages and applications like search engines and automated translators cannot serve these populations.

There has been increasing interest in learning language from more natural input, such as directly from the speech signal, or multi-modal input (e.g. speech and vision). This has several advantages such as removing the need for expensive annotation of speech, being applicable to low resource languages and being more plausible as a model of human language learning.

An important challenge in learning language from spoken input is the fact that the input is not presented in neatly segmented tokens. An auditory signal does not contain neat breaks in between words like the spaces in text. Furthermore, no two realisations of the same spoken word are ever exactly the same. As such, spoken input cannot be represented by conventional word embeddings (e.g. word2vec [4], GloVe [5]). These text-based embeddings are trained to encode word-level semantic knowledge and have become a mainstay in work on sentence representations (e.g. [6, 7]). When we want to learn language directly from speech, we will have to do so in a more end-to-end fashion, without prior lexical level knowledge in terms of both form and semantics.

In previous work [8] we used image-caption retrieval, where given a written caption the model must return the matching image and vice versa. We trained deep neural networks (DNNs) to create sentence embeddings without the use of prior knowledge of lexical semantics (see [7, 9, 10] for other studies on this task). The visually grounded sentence embeddings that arose capture semantic information about the sentence as measured by the Semantic Textual Similarity task (see [11]), performing comparably to text-only methods that require word embeddings.

In the current study we present an image-caption retrieval model that extends our previous work to spoken input. In [12, 13], the authors adapted text based caption-image retrieval (e.g. [9]) and showed that it is possible to perform speech-image retrieval using convolutional neural networks on spectral features. Our work is most closely related to the models presented in [12, 13, 14, 15]. In the current study we improve upon these previous approaches to visual grounding of speech and present state-of-the-art image-caption retrieval results.

The work by [12, 13, 14, 15] and the results presented here are a step towards more cognitively plausible models of language learning as it is more natural to learn language without prior assumptions about the lexical level. For instance, research indicates that the adult lexicon contains many relatively fixed multi-word expressions (e.g., ‘how-are-you-doing’) [16]. Furthermore, early during language acquisition the lexicon consists of entire utterances before a child’s language use becomes more adult-like [16, 17, 18, 19]. Image to spoken-caption retrieval models do not know a priori which constituents of the input are important and have no prior knowledge of lexical level semantics. We probe the resulting model to investigate whether it learns to recognise lexical units in the input without being explicitly trained to do so.

We test two types of acoustic features; Mel Frequency Cepstral Coefficients (MFCCs) and Multilingual Bottleneck (MBN) features. MFCCs are features that can be computed for any speech signal without needing any other data, while the MBN features are ‘learned’ features that result from training a network on top of MFCCs in order to recognise phoneme states. While MBN features have been shown to be useful in several speech recognition tasks (e.g. [20, 21]), learned audio features face the same issue as word embeddings, as humans learn to extract useful features from the audio signal as a result of learning to understand language and not as a separate process. However, the MBN features can still be useful where system performance is more important than cognitive plausibility, for instance in a low resource setting. Furthermore, these features could provide a clue as to what performance would be possible if we had more sophisticated models or more data to improve the feature extraction from the MFCCs in an end-to-end fashion.

In summary, we improve on previous spoken-caption to image retrieval models and investigate whether it learns to recognise words in the speech signal. We show that

our model achieves state-of-the-art results on the Flickr8k database, outperforming previous models by a large margin using both MFCCs and MBN features. We find that our model learns to recognise words in the input signal and show that the deeper layers are better at encoding this information. Recognition performance drops a little in the last two layers as the network abstracts away from the detection of specific words in the input and learns to map the utterances to the joint embedding space. We released the code for this project on github: <https://github.com/DannyMerkx/speech2image/tree/Interspeech19>.

2. Image to spoken-caption retrieval

2.1. Materials

Our model is trained on the Flickr8k database [22]. Flickr8k contains 8,000 images taken from online photo sharing application Flickr.com, for which five English captions per image are available. Annotators were asked to write sentences that describe the depicted scenes, situations, events and entities (people, animals, other objects). Spoken captions for Flickr8k were collected by [12] by having Amazon Mechanical Turk workers pronounce the original written captions. We used the data split provided by [9], with 6,000 images for training and a development and test set both of 1,000 images.

2.2. Image and acoustic features

To extract image features, all images are resized such that the smallest side is 256 pixels while keeping the aspect ratio intact. We take ten 224 by 224 crops of the image: one from each corner, one from the middle and the same five crops for the mirrored image. We use ResNet-152 [23] pretrained on ImageNet to extract visual features from these ten crops and then average the features of the ten crops into a single vector with 2,048 features.

We test two types of acoustic features; Mel Frequency Cepstral Coefficients (MFCCs) and Multilingual Bottleneck (MBN) features. The MFCCs were created using 40 Mel-spaced filterbanks. We use 12 MFCCs and the log energy feature and add the first and second derivatives resulting in 39-dimensional feature vectors. We compute the MFCCs using 25 ms analysis windows with a 5 ms shift.

The MBN features are created using a pre-trained DNN made available by [21]. In short, the network is trained on multilingual speech data (11 languages, no English) to classify phoneme states. The MBN features consist of the outputs of intermediate network layers where the network is compressed from 1500 features to 30 features (see [21] for the full details of the network and training).

2.3. Model architecture

Our multimodal encoder maps images and their corresponding captions to a common embedding space. The idea is to make matching images and captions lie close together and mismatched images and captions lie far apart in the embedding space. Our model consists of two parts; an image encoder and a sentence encoder as depicted in Figure 1. The approach is based on our own text-based model described in [8] and on the speech-based models described in [13, 15] and we refer to those studies for more details. Here, we focus on the differences with previous work.

For the image encoder we use a single-layer linear projec-

Figure 1: *Model architecture: The model consists of two branches with the image encoder on the left and the caption encoder on the right. The audio features consist of n features by t frames and the RNN hidden states by $\mathbf{h}_{t/2}$. Each RNN hidden state has m features which are concatenated for the forward and backward RNN into $2m$ dimensional hidden states. Vectorial attention is applied which weighs and sums the hidden states resulting in the caption embedding. At the top we calculate the cosine similarity between the image and caption embedding ($\mathbf{emb_img}$ and $\mathbf{emb_cap}$).*

tion on top of the pretrained image recognition model, and normalise the result to have unit L2 norm. The image encoder has 2048 input units and 2048 output units.

Our caption encoder consists of three main components. First we apply a 1-dimensional convolutional layer to the acoustic input features. The convolution has a stride of size 2, kernel size 6 and 64 output channels. This is the only layer where the model differs from the text-based model, which features a character embedding layer instead of a convolutional layer. The resulting features are then fed into a bi-directional Gated Recurrent Unit (GRU) followed by a self-attention layer and is lastly normalised to have unit L2 norm.

We use a 3-layer bi-directional GRU which allows the network to capture long-range dependencies in the acoustic signal (see [24] for a more detailed description of the GRU). Furthermore, by making the layer bi-directional we let the network process the output of the convolutional layer from left to right and vice versa, allowing the model to capture dependencies in both directions. We use a GRU with 1024 units, and concatenate the bidirectional representations resulting in hidden states of size 2048. Finally, the self-attention layer computes a weighted sum over all the hidden GRU states:

$$\mathbf{a}_t = \text{softmax}(V \tanh(W\mathbf{h}_t + \mathbf{b}_w) + \mathbf{b}_v) \quad (1)$$

$$\text{Att}(\mathbf{h}_1, \dots, \mathbf{h}_t) = \sum_t \mathbf{a}_t \circ \mathbf{h}_t \quad (2)$$

where $\mathbf{a}_t$ is the attention vector for hidden state $\mathbf{h}_t$ and W , V , $\mathbf{b}_w$, and $\mathbf{b}_v$ indicate the weights and biases. The applied attention is then the sum over the Hadamard product between all hidden states ($\mathbf{h}_1, \dots, \mathbf{h}_t$) and their attention vector. We use 128 units for W and 2048 units for V .

2.4. Training

Following [8], the model is trained to embed the images and captions such that the cosine similarity between image and cap-

tion pairs is larger (by a certain margin) than the similarity between mismatching pairs. This so called hinge loss L as a function of the network parameters θ is given by:

$$L(\theta) = \sum_{(c,i),(c',i') \in B} \left(\max(0, \cos(c, i') - \cos(c, i) + \alpha) + \max(0, \cos(i, c') - \cos(i, c) + \alpha) \right) \quad (3)$$

where $(c, i) \neq (c', i')$. B is a minibatch of correct caption-image pairs (c, i) , where the other caption-image pairs in the batch serve to create mismatched pairs (c, i') and (c', i) . We take the cosine similarity $\cos(x, y)$ and subtract the similarity of the mismatched pairs from the matching pairs such that the loss is only zero when the matching pair is more similar than the mismatched pairs by a margin α . We use importance sampling to select the mismatched pairs; rather than using all the other samples in the mini-batch as mismatched pairs (as done in [8, 15]), we calculate the loss using only the hardest examples (i.e. mismatched pairs with high cosine similarity). While [10] used only the single hardest example in the batch for text-captions, we found that this did not work for the spoken captions. Instead we found that using the hardest 25 percent worked well.

The networks are trained using Adam [25] with a cyclic learning rate schedule based on [26]. The learning rate schedule varies the learning rate smoothly between a minimum and maximum bound which were set to 10^{-6} and 2×10^{-4} respectively. The learning rate schedule causes the network to visit several local minima during training, allowing us to use snapshot ensembling [27]. By saving the network parameters at each local minimum, we can ensemble the embeddings of multiple networks at no extra cost. We use a margin $\alpha = 0.2$ for the loss function. We train the networks for 32 epochs and take a snapshot for ensembling at every fourth epoch. For ensembling we use the two snapshots with the highest performance on the development data and simply sum their embeddings.

The main differences with the approaches described in [13, 15] are the use of multi-layered GRUs, importance sampling, the cyclic learning rate, snapshot ensembling and the use of vectorial rather than scalar attention.

3. Word presence detection

While our model is not explicitly trained to recognise words or segment the speech signal, previous work has shown that such information can be extracted by visual grounding models [15, 28]. [15] use a binary decision task: given a word and a sentence embedding, decide if the word occurs in the sentence. Our approach is similar to the spoken-bag-of-words prediction task described in [28]. Given a sentence embedding created by our model, a classifier has to decide which of the words in its vocabulary occur in the sentence.

Based on the original written captions, our database contains 7,374 unique words with a combined occurrence frequency of 324,480. From these we select words that occur between 50 and a 1,000 times and are over 3 characters long so that there are enough examples in the data that the model might actually learn to recognise them, and to filter out punctuation, spelling mistakes, numerals and most function words. This leaves 460 unique words, mostly verbs and nouns, with a combined occurrence frequency of 87,020 in our data. We construct a vector for each sentence in Flickr8k indicating which of

Table 1: *Image-Caption retrieval results on the Flickr8k test set. R@N is the percentage of items for which the correct image or caption was retrieved in the top N (higher is better). Med r is the median rank of the correct image or caption (lower is better). We also report the 95 percent confidence interval for the R@N scores.*

Model		Caption to Image			
		R@1	R@5	R@10	med r
[12]		-	-	17.9±1.1	-
[15]		5.5±0.6	16.3±1.0	25.3±1.2	48
MFCC-GRU		8.4±0.8	25.7±1.2	37.6±1.3	21
MBN-GRU		12.7±0.9	34.9±1.3	48.5±1.4	11
Char-GRU [8]		27.5±1.2	58.2±1.4	70.5±1.3	4
Model		Image to Caption			
		R@1	R@5	R@10	med r
[12]		-	-	24.3±2.7	-
MFCC-GRU		12.2±2.0	31.9±2.9	45.2±3.1	13
MBN-GRU		16.0±2.5	42.8±3.1	56.1±3.0	8
Char-GRU [8]		38.5±3.0	68.9±2.9	79.3±2.5	2

these words is present. We do not encode multiple occurrences of the same word in one sentence.

The words described above are used as targets for a neural network classifier consisting of a single feed forward layer with 460 units. This layer simply takes an embedding vector as input and maps it to the 460 target words. We then apply the standard logistic function and calculate the Binary Cross Entropy loss to train the network.

We train five word detection networks for both the MFCC and the MBN based caption encoders, in order to see how word presence is encoded in the different neural network layers. We train networks for the final output layer, the three intermediate layers of the GRU and the acoustic features. For the final layer we simply use the output embedding as input to the word detection network. We apply some post-processing to the acoustic features and the intermediate layer outputs to ensure that our word detection inputs are all of the same size. As the intermediate GRU layers produce 2048 features for each time step in the signal, we use average-pooling along the temporal dimension to create a single input vector and normalise the result to have unit L2 norm. The acoustic features consist of 30 (MBN) or 39 (MFCC) features for each time step, so we apply the convolutional layer followed by an untrained GRU layer to the input features, use average-pooling and normalise the result to have unit L2 norm.

The word detection networks are trained for 32 epochs using Adam [25] with a constant learning rate of 0.001. We use the same data split that was used for training the multi-modal encoder, so that we test word presence detection on data that was not seen by either the encoder or the decoder.

4. Results

Table 1 shows the performance of our models on the image-caption retrieval task. The caption embeddings are ranked by cosine distance to the image and vice versa where R@N is the percentage of test items for which the correct image or caption was in the top N results. We compare our models to [12] and [15], and include our own character-based model for comparison. [12] is a convolutional approach, whereas [15] is an approach using recurrent highway networks with scalar atten-

Figure 2: Plots of the F1 scores for the word presence classifiers at 20 equally spaced activation thresholds. The top figure shows the classifiers trained on the MBN model, and the bottom figure the MFCC model.

tion. The character-based model is similar to the model we use here and was trained on the original Flickr8k text captions (see [8] for a full description). Both our MFCC and MBN based model significantly outperform previous spoken caption-to-image methods on the Flickr8k dataset. The largest improvement is the MBN model which outperforms the results reported in [15] by as much as 23.2 percentage points on R@10. The MFCC model also improves on previous results but scores significantly lower than the MBN model across the board, improving as much as 12.3 percentage points over previous work. There is a large performance gap between the text-caption to image retrieval results and the spoken-caption to image results, showing there is still a lot of room for improvement.

The results of the word presence detection task are shown in Figure 2 and Table 2. Figure 2 shows the F1 score for all the classifiers at 20 equally spaced detection thresholds (i.e. a word is classified as ‘present’ if the word detection output is above this threshold). Table 2 displays the area under the curve for the receiver operating characteristic. Even though the MBN model outperforms the MFCC model for all layers we see the same pattern emerging from both the F1 score and the AUC. The performance on the feature level is not much better than random. Predicting ‘not present’ for every word would be the best random guess as this is a heavy majority class in this task. Inspection of the predictions shows that the classifier is indeed heavily biased towards the majority class for the input features. Then we see the performance increasing for the first layer and peaking at the second layer. The performance then drops slightly for the third layer and the attention layer.

Table 2: Area under the curve of the receiver operating characteristic for both models.

Model	AUC				
	input	layer 1	layer 2	layer 3	attention
MBN	.57	.80	.86	.85	.82
MFCC	.54	.68	.80	.75	.75

5. Discussion and Conclusion

We trained an image-caption retrieval model on spoken input and investigated whether it learns to recognise linguistic units in the input. As improvements over previous work we used a 3-layer GRU and employed importance sampling, cyclic learning rates, ensembling and vectorial self-attention. Our results on both MBN and MFCC features are significantly higher than the previous state-of-the-art. The largest improvement comes from using the learned MBN features but our approach also improves results for MFCCs, which are the same features as were used in [15]. The learned MBN features provide better performance whereas the MFCCs are more cognitively plausible input features.

The probing task shows that the model learns to recognise these words in the input. The system is not explicitly optimised to do so, but our results show that the lower layers learn to recognise this form related information from the input. After layer 2, the performance starts to decrease slightly which might indicate that these layers learn a more task-specific representation and it is to be expected that the final attention layer specialises in mapping from audio features to the multi-modal embedding space.

In conclusion, we presented what are, to the best of our knowledge, the best results on spoken-caption to image retrieval. Our results improve significantly over previous approaches for both untrained and trained audio features. In a probing task, we show that the model learns to recognise words in the input speech signal.

We are currently collecting the Semantic Textual Similarity (STS) database in spoken format and the next step will be to investigate whether the model presented here also learns to capture sentence level semantic information and understand language in a deeper sense than recognising word presence. The work presented in [15] has made the first efforts in this regard and we aim to extend this to a larger database with sentences from multiple domains. Furthermore, we want to investigate the linguistic units that our model learns to recognise. In the current study, we only investigated whether the model learns to recognise words, but the potential benefit of our model is that it might learn multi-word statements or might even learn to look at sub-lexical level information. [14, 29] have recently shown that the speech-to-image retrieval approach can be used to detect word boundaries and even discover sub-word units. Our interest is in investigating how these word and sub-word units develop over training and through the network layers.

6. Acknowledgements

The research presented here was funded by the Netherlands Organisation for Scientific Research (NWO) Gravitation Grant 024.001.006 to the Language in Interaction Consortium. This work was carried out on the Dutch national e-infrastructure with the support of SURF Cooperative.

7. References

- [1] X. Wang, H. Pham, P. Yin, and G. Neubig, “A tree-based decoder for neural machine translation,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, November 2018.
- [2] R. Kiros, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler, “Skip-thought vectors,” in *Advances in Neural Information Processing Systems* 28, 2015, pp. 3294–3302.
- [3] Z.-Q. Wang and D. Wang, “A joint training framework for robust automatic speech recognition,” *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 24, no. 4, pp. 796–806, 2016.
- [4] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Advances in Neural Information Processing Systems* 26, 2013, pp. 3111–3119.
- [5] J. Pennington, R. Socher, and C. Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2014, pp. 1532–1543.
- [6] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, “Supervised Learning of Universal Sentence Representations from Natural Language Inference Data,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, September 2017, pp. 670–680.
- [7] D. Kiela, A. Conneau, A. Jabri, and M. Nickel, “Learning visually grounded sentence representations,” in *Proceedings of NAACL-HLT 2018*. Association for Computational Linguistics, June 2018, pp. 408–418.
- [8] D. Merx and S. L. Frank, “Learning semantic sentence representations from visually grounded language without lexical knowledge,” *Natural Language Engineering*, in press.
- [9] A. Karpathy and L. Fei-Fei, “Deep visual-semantic alignments for generating image descriptions,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [10] F. Faghri, D. J. Fleet, R. Kiros, and S. Fidler, “VSE++: improved visual-semantic embeddings,” *CoRR*, vol. abs/1707.05612, 2017.
- [11] E. Agirre, C. Banea, D. Cer, M. Diab, A. Gonzalez-Agirre, R. Mihalcea, G. Rigau, and J. Wiebe, “Semeval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation,” in *SemEval*. ACL, 2016, pp. 497–511.
- [12] D. Harwath and J. Glass, “Deep multimodal semantic embeddings for speech and images,” in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 237–244.
- [13] D. Harwath, A. Torralba, and J. Glass, “Unsupervised learning of spoken language with visual context,” in *Advances in Neural Information Processing Systems* 29, 2016, pp. 1858–1866.
- [14] D. Harwath, A. Recasens, D. Surís, G. Chuang, A. Torralba, and J. Glass, “Jointly discovering visual objects and spoken words from raw sensory input,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11210 LNCS, pp. 659–677, 2018.
- [15] G. Chrupała, L. Gelderloos, and A. Alishahi, “Representations of language in a model of visually grounded speech signal,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, July 2017, pp. 613–622.
- [16] M. Tomasello, “First steps toward a usage-based theory of language acquisition,” *Cognitive Linguistics*, vol. 11, no. 1/2, pp. 61–82, 2000.
- [17] M. D. S. Braine and M. Bowerman, “Children’s first word combinations,” *Monographs of the Society for Research in Child Development*, vol. 41, no. 1, pp. 1–104, 1976.
- [18] J. M. Pine and E. Lieven, “Reanalysing rote-learned phrases: individual differences in the transition to multi-word speech,” *Journal of Child Language*, vol. 20, p. 551571, 1993.
- [19] E. Lieven, H. Behrens, J. Speares, and M. Tomasello, “Early syntactic creativity: a usage-based approach,” *Journal of Child Language*, vol. 30, no. 2, pp. 333–370, 2003.
- [20] Q. B. Nguyen, J. Gehring, M. Müller, S. Stücker, and A. Waibel, “Multilingual shifting deep bottleneck features for low-resource asr,” in *2014 IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 5607–5611.
- [21] R. Fer, P. Matejka, F. Grezl, O. Plchot, K. Vesely, and J. H. Cernocky, “Multilingually trained bottleneck features in spoken language recognition,” *Computer Speech & Language*, vol. 46, no. Supplement C, pp. 252–267, 2017.
- [22] M. Hodosh, P. Young, and J. Hockenmaier, “Framing image description as a ranking task: Data, models and evaluation metrics,” *Journal of Artificial Intelligence Research*, vol. 47, no. 1, pp. 853–899, 2013.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [24] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, “Empirical evaluation of gated recurrent neural networks on sequence modeling,” in *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014.
- [25] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of the International Conference on Learning Representations (ICLR) 2015*, May 2015.
- [26] L. N. Smith, “Cyclical learning rates for training neural networks,” in *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, March 2017, pp. 464–472.
- [27] G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger, “Snapshot Ensembles: Train 1, get M for free,” in *International Conference on Learning Representations (ICLR) 2017*, April 2017.
- [28] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 2017, pp. 3677–3681, August 2017.
- [29] J. Drexler and J. Glass, “Analysis of audio-visual features for unsupervised speech recognition,” in *2017 International Workshop on Grounding Language Understanding, GLU*, August 2017, pp. 57–61.