

VISUAL FEATURES FOR CONTEXT-AWARE SPEECH RECOGNITION

Abhinav Gupta, Yajie Miao, Leonardo Neves, and Florian Metze

Language Technologies Institute, Carnegie Mellon University
Pittsburgh, PA; U.S.A.

abhinavgupta94@gmail.com, {ymiao|lneves|fmetze}@cs.cmu.edu

ABSTRACT

Automatic transcriptions of consumer generated multi-media content such as “Youtube” videos still exhibit high word error rates. Such data typically occupies a very broad domain, has been recorded in challenging conditions, with cheap hardware and a focus on the visual modality, and may have been post-processed or edited.

In this paper, we extend our earlier work on adapting the acoustic model of a DNN-based speech recognition system to an RNN language model, and show how both can be adapted to the objects and scenes that can be automatically detected in the video. We are working on a corpus of “how-to” videos from the web, and the idea is that an object that can be seen (“car”), or a scene that is being detected (“kitchen”) can be used to condition both models on the “context” of the recording, thereby reducing perplexity and improving transcription. We achieve good improvements in both cases, and compare and analyze the respective reductions in word error rate.

We expect that our results can be useful for any type of speech processing in which “context” information is available, for example in robotics, man machine interaction, or when indexing large audio-visual archives, and should ultimately help to bring together the “video-to-text” and “speech-to-text” communities.

Index Terms— audio-visual speech recognition, multi-modal processing, deep learning

1. INTRODUCTION

Robustness or adaptation to signal variability is a key challenge if automatic speech recognition (ASR) systems are to become universally useful. One way in which this could be achieved is to adapt both the acoustic model and the language model to the broad “context” of the input. By “context”, we mean essentially anything that is known about the input speech.

State-of-the-art recognition accuracy on a wide range of acoustic modeling tasks is defined by DNNs [1, 2, 3], or variants thereof. On consumer-generated content (“Youtube videos”) however, even DNN models exhibit word error rates (WERs) above 40% [4], although no standardized test set exists. Other work reports significantly lower [5] or higher WERs [6], showing the wide variability that exists in such data. Most recently, low word error rates have been recorded in an extremely high resource setup [7].

An effective strategy to deal with variability is to incorporate additional, longer-term knowledge explicitly into DNN models: [8, 9, 10, 11, 12] study the incorporation of speaker-level i-vectors to smooth out the effect of speaker variability. Time Delay Neural Networks [13, 14] use wide temporal input windows to improve robustness dynamically. [15] extracts long-term averages from the audio signal to adapt a DNN acoustic model. Similarly, in [16], we learn a DNN-based extractor to model the speaker-microphone distance information dynamically on the frame level. Then distance-aware DNNs are built by appending these descriptors to the DNN inputs.

It is an important distinction that our work *does not require* localization of lip regions and/ or extraction of frame-synchronous visual features (lip contours, mouth shape, SIFT, landmarks, etc.), as is the case in “traditional” audio-visual ASR [17, 18, 19, 20, 21], which has been developed mostly with a focus on noise robustness. For the majority of our data, lip-related information is not available at all, or the quality is extremely poor.

In this paper, we present an extension of and comparison to previously published work [5], by adapting not only the acoustic model, but also the language model of an ASR system to the visual “context” that is present in the video stream of open-domain Internet video. Our approach is based on deep learning, and involves two major steps:

First, we extract visual features using deep Convolutional Neural Networks (CNNs) trained for object recognition and scene labeling tasks. We extract such information from a single, random frame within an utterance only, rather than at the level of each frame, but other levels of granularity, or smoothing approaches are also possible. Thus, we do not require perfect alignment between audio and video channels, which is often almost impossible to achieve on data that has been

This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575.

Submitted for publication in fond memory of Yajie Miao († 2016).

collected “in the wild”. Our “context vector” is therefore an n -dimensional representation of the visuals which are present while an utterance is being spoken.

Then, we adapt the acoustic model of the recognizer using a framework in which the residual error at the feature inputs of a DNN is reduced with an adaptation network. This network is trained on the context vector, and predicts a linear shift of the main DNN’s input features. The central idea is somewhat similar to ResNets [22], and was originally developed for i-vector based adaptation [11]. It works well for adaptation to other knowledge sources as well.

For first-pass decoding, we use an in-domain 3-gram language model. To adapt the language model of the recognizer, we then re-rank 30-best lists with an RNN language model, which has been conditioned on the same segment-level “context vector” as the acoustic model. We show that this approach results in significant reductions in perplexity and also reduces word error rate.

2. EXTRACTION OF VISUAL FEATURES

The extraction of visual features follows our previous work on adaptation of DNNs using speaker attributes [23] and visual features [5].

Suppose we are dealing with an utterance u , which has the acoustic features $\mathbf{O} = \{o_1, o_2, \dots, o_T\}$, where T is the total number of speech frames. On a video transcribing task, there always exists a video segment corresponding to u . This video segment is represented as $\mathbf{V} = \{v_1, v_2, \dots, v_N\}$, where N is the number of video frames. The video frames are sampled normally at a lower sampling rate than the speech frames, i.e. $N < T$. From this segment, we randomly select a frame v_n which serves as the image representation for the utterance. Then two types of image features are extracted from v_n .

2.1. Object Features

Our first type of visual information is derived from object recognition, a task on which deep learning has accomplished tremendous success [24]. The intuition is that object features contain information regarding the acoustic environment. For example, classifying an image to the classes “computer keyboard” and “monitor” indicates that the speech segment has been recorded in an office.

We extract this object information using a deep CNN model which has been trained on a comprehensive object recognition dataset, a 1.2 million image subset of ImageNet [25] used for the 2012 ILSVR challenge, and the resulting CNN model is referred to as OBJECT-CNN. Then, on our target ASR task, the video frame v_n is fed into the CNN model, from which we get the distribution (posterior probabilities) over the object classes. These probabilities encode the object-related information that are finally incorporated into DNN acoustic models.

The OBJECT-CNN network follows the standard AlexNet architecture [25]. The network contains 5 convolution layers which use the rectifier non-linearity (ReLU) [26] as the activation function. In the first and second convolution layers, a local response normalization (LRN) layer is added after the ReLU activation, and a max pooling layer follows the LRN layer. In the third and fourth convolution layers, we do not apply the LRN and pooling layers. In the fifth convolution layer, we only apply the max pooling layer, without LRN being applied. 3 fully-connected (FC) layers are placed on top of the convolution layers. The first and second FC layers have 4096 neurons, whereas the number of neurons in the last FC layer is equal to the number of classes, 1000 in our case. Model training optimizes the standard cross-entropy (CE) objective. The resulting OBJECT-CNN achieves a 20% top-5 error rate on the ILSVRC 2012 testing set.

2.2. Place Features

The utility of the object features comes from the “place” information that is implicitly encoded by the object classification results. It is then natural to utilize place features in a more explicit way. To achieve this, we train a deep CNN model meant for the scene labeling task. Given a video frame, the classification outputs from this PLACE-CNN encode the place information, which is then incorporated into acoustic models. For convenience of formulation, the resulting visual feature vector for this utterance u is represented as f_u .

In order to extract place information, we train the PLACE-CNN network on the MIT Places dataset [27] which contains 2.5 million images belonging to 205 scene categories. Examples of the scenes include “dining room”, “coast”, “conference center”, “courtyard”, etc. We use the complete set of 2.5 million images for training, and follow the same image pre-processing as used on ImageNet (Section 2.1). The architecture of the PLACE-CNN is almost the same as that of the OBJECT-CNN. The only difference is that in the final FC layer, the PLACE-CNN has 205 neurons corresponding to the 205 scene classes, whereas the OBJECT-CNN contains 1000 neurons.

3. EXPERIMENTAL SETUP

We chose to investigate context-aware ASR on a dataset of real-world English instructional videos, which we had downloaded from online video archives [16, 23]. These videos have been uploaded by social media users to share expertise on specific tasks (e.g., oil change, sandwich making, etc.). ASR on these videos is challenging because they have been recorded in various environments (e.g., office, kitchen, baseball field, train, etc.), giving us a variety of contexts, yet they are rich in speech, making them suitable for the proposed work. Our main training set comprises 90 hours of speech (3900 videos), and we use 4 hours for testing (156 videos).

We used Kaldi [28] and PDNN [29] for our experiments, training a 5-layer DNN acoustic model using cross-entropy [5]. For decoding, a trigram language model (LM) is trained on the training transcripts. This LM is then interpolated with another trigram LM trained on an additional set of 270 hours transcriptions of instructional videos. The complete set of 360 hours is also used for training the RNN language model.

4. ACOUSTIC MODEL ADAPTATION

In previous work [11, 12], we presented a framework to perform speaker adaptive training (SAT) for DNN models. This approach requires an i-vector [30] to be extracted for each speaker. Based on the well-trained speaker-independent (SI) DNN, a separate adaptation neural network is learned to convert i-vectors into speaker-specific linear feature shifts. Adding these shifts to the original DNN inputs produces a speaker-normalized feature space. Parameters of the SI-DNN are re-updated in this new space, generating the SAT-DNN model. This framework has also been applied successfully to descriptors of speaker-microphone distance [16], and we find it to be more robust than straightforward feature concatenation [12].

We port this idea to visual input features, which enables us to conduct “context” adaptation for DNNs, simply replacing the i-vector representation with the visual features. An adaptation network is learned to take the visual features as inputs and generate an adaptive feature space with respect to the visual descriptors. Note that in this case, the linear feature shifts generated by the adaptation network are utterance-specific rather than speaker-specific. Re-updating the parameters of the DNN in the normalized feature space gives us the adaptively trained “video adaptive training” VAT-DNN model [5]. This VAT-DNN model takes advantage of the visual features as additional knowledge, and generalizes better to unseen variability. In our setup, we generate 100-dimensional utterance-level visual “context” features by projecting the output vector (after soft-max) of the PLACE-CNN and OBJECT-CNN (either individually, or in concatenation) down to 100 dimensions using Principal Component Analysis (PCA). The outputs of the adaptation network are 40-dimensional shifts to IMEL features.

5. LANGUAGE MODEL ADAPTATION

To adapt the language model, we used the same features that we also use for adapting the acoustic model (100-dimensional PCA projections of place and scene) as “topic” information in a context dependent Recurrent Neural Network (RNN) language model [31]. The vocabulary contains about 35k words. A two-layer bidirectional LSTM with an embedding layer size of 900 and 1024 cells per layer gave lowest perplexities in initial cross-validation experiments on 360 hours of data. This architecture was thus adopted for the experiments below. We

Base-line	Object F.	Place F.	O. + P. F.	i-vectors	All Features
23.4%	22.5%	22.5%	22.3%	22.0%	21.5%

Table 1: Word error rates when applying acoustic model adaptation using object features, place features, a combination thereof [5], i-vectors [5], and a combination of the two visual features with i-vectors (“all”).

used tanh non-linearities and 0.5 as the dropout factor, without any additional regularization except gradient clipping (at 100). The initial learning rate was 0.01; training used Ada-grad [32] and a batch size of 128. The network was implemented in Lasagne [33].

We provide the adaptation vector only at the beginning of the sentence, although it might make sense to provide it also at intermediate steps, as the average sentence length is 18 words.

6. EXPERIMENTS

To reduce the dimensionality of the adaptation feature and to facilitate comparison with earlier work on i-vector adaptation, we reduce the dimensionality of the place and object features to 100 (from 1000 and 205) using PCA, estimated on the training part only of the audio-visual dataset.

Table 1 shows the result of adapting the DNN acoustic model with visual features, and i-vectors for comparison, as well as a combination of visual features and i-vectors. Gains are consistent, and quite complementary when using the concatenation of visual features and i-vectors for adaptation. Also, in all cases, the adaptation network method outperforms simply concatenating the adaptation vector to the input features.

Next, we use the same method to adapt the language model to the visual information. To find the best meta parameters for the LSTM language model, we performed 5-fold cross-validation on the entire 360 hours of training data, and averaged the results. Figure 1 shows that conditioning the LSTM LM on video features reduces perplexity from 89 to 74 on training data, which is a significant reduction, which we find does also carry over to the word error rates on unseen test data.

We generated 30-best lists using the baseline acoustic models (with a WER of 23.4%), which had an oracle WER of 15.6%, and re-ranked them with all 5 neural network language models (NN LMs), averaging the language model scores. Using the concatenation of object and place features as inputs to the NN LM, we achieve a word error rate of 22.6%, which is very close to the performance achieved with the adaptation of the acoustic model.




Fig. 1: Training and validation perplexity of the NN LMs trained with and without visual features, averaged across 5 folds, on 360 hours of data. Parameters were optimized for the adaptive LMs, hence the baseline LM converges quickly. Validation perplexity is lower than training perplexity initially because it is measured at the end of each training iteration, while training perplexity is being computed while processing the data and updating the model.




Fig. 2: Keyframes for two typical videos in our dataset. The baseline WER is 27.6% for the “home” video on the left, and 47.7% for the “train” video on the right. Acoustic model adaptation does not improve the “home” video at all, but reduces WER to 38.2% on the “train” video. Language model adaptation improves WER on both videos slightly, to 26.6% and 43.2% respectively.

7. ANALYSIS OF RESULTS

For both acoustic and language model adaptation, we performed some more in-depth analysis to see where gains are mostly coming from.

We manually inspected those videos on which we observed more than 10% relative WER reductions, and find that they have been recorded either in outdoor environments (e.g. baseball field, airport apron, street, etc.), or in non-typical indoor conditions (e.g. kitchen, music studio, etc.), where music/ noise may interfere with the actual speech a lot. Adding the scene descriptors helps the DNN model normalize the acoustic characteristics of these rare testing conditions, and thus benefits the generalization to unseen testing speech. We then labeled all 156 testing videos as either “typical indoors” (e.g. office) and “other” (noisy indoors, outdoors), and analyzed the relative improvements with the of a system adapted with PLACE-CNN features only, and find that the “quiet

indoors” videos get improved from 22.1% WER to 21.7%, while “other” videos get improved from 27.6% to 25.7%. “Other” videos thus get improved by 7% relative, while clean videos get improved by 2% only.

When training the NN LM on 90 hours of data only, adaptation with OBJECT-CNN features results in a perplexity of 94.7, while adaptation with PLACE-CNN features gives a perplexity of 98.9. It seems intuitive that “objects” would be slightly more salient for the topic of a “how-to” video than the scene.

Figure 2 shows typical keyframes from our database, and the typical pattern of improvements: acoustic model adaptation tends to give significant improvements on “outdoor” videos only (55 videos), while language model adaptation tends to give smaller improvements across the board (only 30 of 156 videos deteriorate).

8. CONCLUSION AND FUTURE WORK

In this paper, we described a system that extracts context information that is relevant for speech processing from the visual channel of the video. We showed that the information can be incorporated in both acoustic *and* language models, and that this approach leads to systematic and consistent improvements. These observations are in line with recent work on multi-modal machine translation [34].

We are currently expanding the acoustic model adaptation experiments to the larger (360 hours) version of the corpus, and expect to see further performance improvements by combining both acoustic and language model adaptation. We are also experimenting with different and better ways of incorporating the video features into the language model, and attempt more insightful analysis of the results, e.g. how much do different types of features contribute to the different models (e.g., do the scene features contribute relatively more to the acoustics, while the object features contribute more to the language model?), and what types of errors are being reduced (e.g., nouns? verbs? semantically confusing errors?).

In the long term, this work should help to improve fully end-to-end “video-to-text” approaches, which generate image or video “summaries” based on multi-modal embeddings, and reference “captions” [35, 36, 37], rather than speech recognition transcriptions.

9. REFERENCES

- [1] G. E. Dahl, D. Yu, L. Deng, and A. Acero, “Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 1, pp. 30–42, 2012.
- [2] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al., “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, 2012.

- [3] F. Seide, G. Li, X. Chen, and D. Yu, "Feature engineering in context-dependent deep neural networks for conversational speech transcription," in *Proc. ASRU*. IEEE, 2011, pp. 24–29.
- [4] H. Liao, E. McDermott, and A. Senior, "Large scale deep neural network acoustic modeling with semi-supervised training data for youtube video transcription," in *Proc. ASRU*. IEEE, 2013, pp. 368–373.
- [5] Y. Miao and F. Metze, "Open-domain audio-visual speech recognition: A deep learning approach," in *Proc. INTERSPEECH*, San Francisco, CA; U.S.A., Sept. 2016, ISCA.
- [6] F. Metze, S. Rawat, and Y. Wang, "Improved audio features for large-scale multimedia event detection," in *Proc. ICME*, Chengdu; China, July 2014, IEEE.
- [7] H. Soltau, H. Liao, and H. Sak, "Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition," *ArXiv e-prints*, 2016, <https://arxiv.org/abs/1610.09975>.
- [8] G. Saon, H. Soltau, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *Proc. ASRU*. 2013, pp. 55–59, IEEE.
- [9] V. Gupta, P. Kenny, P. Ouellet, and T. Stafylakis, "I-vector-based speaker adaptation of deep neural networks for french broadcast audio transcription," in *Proc. ICASSP*. IEEE, 2014.
- [10] A. Senior and I. Lopez-Moreno, "Improving dnn speaker independence with i-vector inputs," in *Proc. ICASSP*. IEEE, 2014.
- [11] Y. Miao, H. Zhang, and F. Metze, "Towards speaker adaptive training of deep neural network acoustic models," in *Proc. INTERSPEECH*, Singapore, Sept. 2014, ISCA.
- [12] Y. Miao, H. Zhang, and F. Metze, "Speaker adaptive training of deep neural network acoustic models using i-vectors," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 23, no. 11, pp. 1938–1949, Nov. 2015.
- [13] A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, and K. J. Lang, "Phoneme recognition using time-delay neural networks," *IEEE transactions on acoustics, speech, and signal processing*, vol. 37, no. 3, pp. 328–339, 1989.
- [14] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *Proceedings of INTERSPEECH*. ISCA, 2015, pp. 2440–2444.
- [15] K. Vesel, S. Watanabe, M. Karafi, J. Honza, et al., "Sequence summarizing neural network for speaker adaptation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5315–5319.
- [16] Y. Miao and F. Metze, "Distance-aware DNNs for robust speech recognition," in *Proc. INTERSPEECH*, Dresden, Germany, Sept. 2015, ISCA.
- [17] S. Dupont and J. Luetttin, "Audio-visual speech modeling for continuous speech recognition," *IEEE Transactions on Multimedia*, vol. 2, no. 3, pp. 141–151, 2000.
- [18] T. Chen, "Audiovisual speech processing," *IEEE Signal Processing Magazine*, vol. 18, no. 1, pp. 9–21, 2001.
- [19] C. Bregler, H. Hild, S. Manke, and A. Waibel, "Improving connected letter recognition by lipreading," in *Proc. ICASSP*. IEEE, 1993, vol. 1, pp. 557–560.
- [20] G. Potamianos, J. Luetttin, and C. Neti, "Hierarchical discriminant features for audio-visual lvesr," in *Proc. ICASSP*. IEEE, 2001, vol. 1, pp. 165–168.
- [21] B. P. Yuhas, M. H. Goldstein, and T. J. Sejnowski, "Integration of acoustic and visual speech signals using neural networks," *IEEE Communications Mag.*, vol. 27, no. 11, pp. 65–71, 1989.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.
- [23] Y. Miao, L. Jiang, H. Zhang, and F. Metze, "Improvements to speaker adaptive training of deep neural networks," in *Proc. IEEE Workshop on Spoken Language Technology*, South Lake Tahoe, NV; USA, Dec. 2014, IEEE, Best poster presentation.
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems (NIPS)*, 2012, pp. 1097–1105.
- [25] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009, pp. 248–255.
- [26] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *AISTATS*, 2011, vol. 15, p. 275.
- [27] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Advances in neural information processing systems (NIPS)*, 2014, pp. 487–495.
- [28] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, and J. Silovsky, "The kaldi speech recognition toolkit," in *Proc. ASRU*. IEEE, 2011.
- [29] Y. Miao, "Kaldi+pdnn: building dnn-based asr systems with kaldi and pdnn," *arXiv preprint arXiv:1401.6984*, 2014.
- [30] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.
- [31] T. Mikolov and G. Zweig, "Context dependent recurrent neural network language model," in *Proc. SLT*. IEEE, 2012, pp. 234–239.
- [32] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Res.*, vol. 12, pp. 2121–2159, July 2011.
- [33] S. Dieleman, J. Schlüter, C. Raffel, E. Olson, S. K. Sønderby, D. Nouri, D. Maturana, M. Thoma, E. Battenberg, J. Kelly, et al., "Lasagne: First release," *Zenodo*, 2015.
- [34] D. Elliott, S. Frank, K. Sima'an, and L. Specia, "Multi30K: Multilingual English-German Image Descriptions," *CoRR*, vol. abs/1605.00459, 2016.
- [35] R. Kiros, R. Salakhutdinov, and R. S. Zemel, "Multimodal neural language models," in *Proc. ICML*, June 2014.
- [36] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko, "Translating videos to natural language using deep recurrent neural networks," *arXiv preprint arXiv:1412.4729*, 2014.
- [37] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and tell: A neural image caption generator," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3156–3164.