

On the Contributions of Visual and Textual Supervision in Low-resource Semantic Speech Retrieval

Ankita Pasad¹, Bowen Shi¹, Herman Kamper², Karen Livescu¹

¹Toyota Technological Institute at Chicago, USA

²Dept. E&E Engineering, Stellenbosch University, South Africa

{ankitap,bshi,klivescu}@ttic.edu, kamperh@sun.ac.za

Abstract

Recent work has shown that speech paired with images can be used to learn semantically meaningful speech representations even without any textual supervision. In real-world low-resource settings, however, we often have access to some transcribed speech. We study whether and how visual grounding is useful in the presence of varying amounts of textual supervision. In particular, we consider the task of semantic speech retrieval in a low-resource setting. We use a previously studied data set and task, where models are trained on images with spoken captions and evaluated on human judgments of semantic relevance. We propose a multitask learning approach to leverage both visual and textual modalities, with visual supervision in the form of keyword probabilities from an external tagger. We find that visual grounding is helpful even in the presence of textual supervision, and we analyze this effect over a range of sizes of transcribed data sets. With ~ 5 hours of transcribed speech, we obtain 23% higher average precision when also using visual supervision.

Index Terms: speech search, multi-modal modelling, visual grounding, semantic retrieval, multitask learning

1. Introduction

In many languages and domains, there can be insufficient data to train modern large-scale models for common speech processing tasks. Recent work has begun exploring alternative sources of weak supervision, such as visual grounding from images [1–6]. Such work has developed approaches for training models for tasks such as cross-modal retrieval [2, 3]; unsupervised learning of word-like and phone-like units [7–10]; and retrieval tasks like keyword search, query-by-example search, and semantic search [6, 11, 12]. Much of this recent work is in the context of *zero* textual supervision, that is, using no transcribed speech, and the results show that visual grounding alone provides a strong supervisory signal. However, in many low-resource settings we also have access to a small amount of textual supervision. This is the setting that we address here. In particular, we consider the task of semantic speech retrieval using models trained on a data set of images paired with spoken and written captions.

In semantic speech retrieval, the input is a textual query, a written word in our setting, and the task is to find utterances in a speech corpus that are semantically relevant to the query [13, 14]. The query word need not exactly occur in the retrieved utterances; for example, the query *beach* should retrieve the exactly matched utterance “a dog retrieves a branch from a *beach*” as well as the semantically matched utterance “people at an ocean-side resort.” One common approach is to cascade an automatic speech recognition (ASR) system with a text-based information retrieval method [15]. High-quality ASR, however, requires significant amounts of transcribed speech audio and text for

language modelling. Kamper *et al.* [6, 11] proposed learning a semantic speech retrieval model solely from images of natural scenes paired with unlabelled spoken captions. This approach uses an external visual tagger to automatically obtain text labels from the training images, and these are used as targets for training a speech-to-keyword network. No textual supervision is thus used. For analysis, the approach was compared to a supervised model trained only with text supervision. The visually supervised model was shown to have an edge over the textually trained model in retrieving non-exact keyword matches. In retrieving exact semantic matches, the text-based model proved superior.

Based on these observations, we hypothesize that visual- and text-based supervision can be complementary for semantic speech retrieval in low-resource settings where both are available. Visual supervision could, e.g., provide a signal to distinguish acoustically similar but semantically unrelated words. Here we consider a regime where the amount of transcribed speech audio is not enough to train a full ASR system, but is enough to provide a useful additional supervision signal. Using a corpus of images with spoken captions, we propose a multitask learning (MTL) approach where visual supervision is obtained from an external image tagger and text supervision is obtained from (a limited amount of) transcribed spoken captions. Each type of target is given as a bag of words. Additionally, we include a representation learning loss to encourage similarity between the visual and spoken representations.

We consider different amounts of transcribed speech audio to investigate the effect of MTL in different low-resource settings and to understand the trade-off between visual- and textual supervision. Experiments using human semantic relevance judgments show that our new MTL approach incorporating both visual and text information performs better than when using either one of the supervision modalities in isolation. We achieve consistent improvements over the entire range of transcribed data set sizes considered, with the benefits of visual supervision being largest when only small amounts of transcriptions are available. This result demonstrates that the benefit of visual grounding remains even when textual supervision is available.

2. Approach

As a starting point, we use the model of [6], which addresses the semantic retrieval task using a training set of images and their spoken captions with no textual supervision at all. In this approach, a *vision model*—an image tagger trained on an external data set of tagged images—is used to produce a bag of semantic keywords for each image, along with their posterior probabilities. These posteriors serve as supervision for a *speech model* that takes as input a spoken utterance and outputs a bag of keywords describing the utterance. At test time, these predicted keywords are used for semantic speech retrieval, by retrieving those utter-

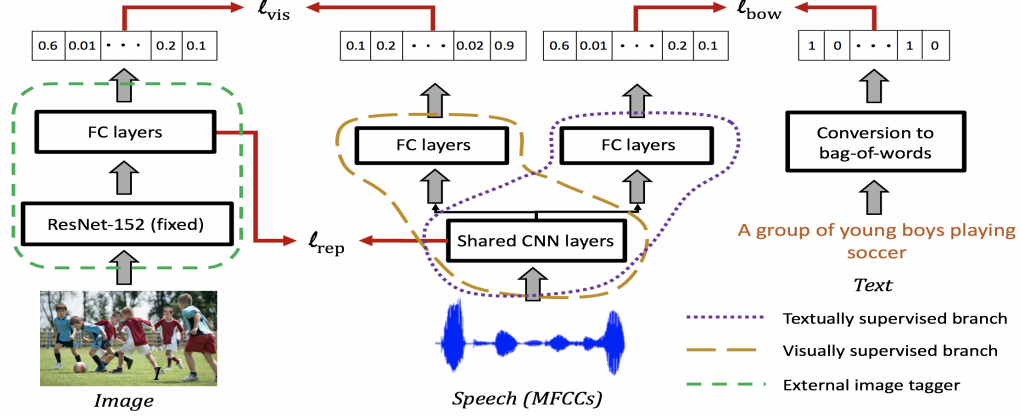


Figure 1: The multitask speech model with visual supervision (right) and textual supervision (left). Either of the textually supervised and visually supervised branches can be used during inference.

ances for which the model predicts a probability higher than a given *detection threshold* for an input query word.

We make two key changes to the model of [6], as illustrated in Figure 1. First, we consider the case where we have both the visual supervision (the vector of keyword probabilities) and textual supervision, obtained from transcriptions of the training utterances. Specifically, for the textual supervision, we convert each transcription into a bag of ground-truth content words occurring in the utterance. The textual and visual supervision are used together in a multitask learning (MTL) approach, where our model consists of two branches, one that produces the “visual keywords” and one that produces “textual keywords”.

The second change we make is to add a *representation loss* term. The purpose of this is to encourage intermediate representations learned by the speech model to match those in the external visual tagger. This loss is similar to ones used in recent work on unsupervised joint learning of visual and speech representations, e.g. for cross-modal retrieval [2]. In our case the representation loss can be viewed as a regularizer or as a third task in the MTL framework. This loss helps us take advantage of the fact that the visual tagger is trained on much more data than the speech model, so we expect its internal representations themselves to be useful during training of the speech network.

As a final, minor change, we use a stronger pre-trained visual network (ResNet-152) in the visual tagger relative to [6] (more details about the vision model are given in Section 3).

2.1. Model Details

Each training utterance $U = \mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_T$ is paired with an image I . Each frame \mathbf{u}_t is an acoustic feature vector, MFCCs in our case.¹ The vision model provides weak labels, $\mathbf{y}_{\text{vis}} \in [0, 1]^{N_{\text{vis}}}$, where N_{vis} is the number of keywords in the visual tag set. This serves as the ground truth for the visually supervised branch output, $f_{\text{vis}}(U) = \hat{\mathbf{y}}_{\text{vis}}$, of the speech model.

Each training utterance U is also optionally paired with a multi-hot bag-of-words vector $\mathbf{y}_{\text{bow}} \in \{0, 1\}^{N_{\text{bow}}}$, obtained from the transcriptions of the spoken captions, where N_{bow} is the number of unique keywords in the text labels and each dimension indicates the presence of absence of a particular word. This vector serves as a ground truth for the output of the textually supervised branch, $f_{\text{bow}}(U) = \hat{\mathbf{y}}_{\text{bow}}$, of the speech model.

¹Earlier work on this data set compared MFCCs to filterbank features, and found that MFCCs worked similarly or better [11].

These task-specific supervised losses are optimized jointly with an unsupervised multi-view representation loss. In particular, we use a margin-based contrastive loss [16] between the speech representation \mathbf{s} and visual representation \mathbf{v} at an intermediate layer within each model.

The total loss is a weighted sum of these three losses: $\ell = \alpha_{\text{vis}} \cdot \ell_{\text{vis}} + \alpha_{\text{bow}} \cdot \ell_{\text{bow}} + (1 - \alpha_{\text{bow}} - \alpha_{\text{vis}}) \cdot \ell_{\text{rep}}$ where each of the loss terms is a function of a training utterance U and either (1) a corresponding image representation \mathbf{v} , (2) a visual target \mathbf{y}_{vis} , or (3) a textual target \mathbf{y}_{bow} (if available). Both of the supervised task losses, ℓ_{vis} and ℓ_{bow} , are summed cross entropy losses between the predicted and ground-truth vectors, as follows (where $\text{sup} \in \{\text{vis}, \text{bow}\}$)

$$\ell_{\text{sup}} = - \sum_{w=1}^{|N_{\text{sup}}|} \{ y_{\text{sup},w} \log \hat{y}_{\text{sup},w} + (1 - y_{\text{sup},w}) \log [1 - \hat{y}_{\text{sup},w}] \} \quad (1)$$

where \hat{y}_{sup} is a function of the speech input and \mathbf{y}_{sup} is function of the image/transcription. The representation loss is a contrastive loss, similar to ones used in prior work on multi-view representation learning [2, 16, 17], computed by sampling a fixed number of negative examples within a mini-batch (size B) corresponding to both images and utterances:

$$\ell_{\text{rep}} = \left\{ \frac{1}{|V|} \sum_{\mathbf{v}' \in V} \max[0, m + d_{\cos}(\mathbf{v}, \mathbf{s}) - d_{\cos}(\mathbf{v}', \mathbf{s})] + \frac{1}{|S|} \sum_{\mathbf{s}' \in S} \max[0, m + d_{\cos}(\mathbf{v}, \mathbf{s}) - d_{\cos}(\mathbf{v}, \mathbf{s}')] \right\} \quad (2)$$

where $\{\mathbf{v}, \mathbf{s}\}$ are the representations of the correct vision-speech pair; $\{\mathbf{v}', \mathbf{s}\}$ and $\{\mathbf{v}, \mathbf{s}'\}$ are negative (non-matching) pairs; $d_{\cos}(\mathbf{v}, \mathbf{s})$ is the cosine distance between the representations; $n_{\text{neg}} = |V| = |S|$ is the number of negative pairs; and m is a margin indicating the minimum desired difference between the positive-pair distances and negative-pair distances.

3. Experimental setup

3.1. Data

We use three training sets and one evaluation data set, identical to those used in prior work [6], allowing for direct comparison.

(A) *Image-text pairs*: The union of MSCOCO [18] and Flickr30k [19], with $\sim 149\text{k}$ images ($\sim 107\text{k}$ training, $\sim 42\text{k}$ dev) paired with 5 written captions each. This is an external data

set used to train the image tagger before our MTL approach is applied. We are taking advantage of the fact that, in contrast to speech, labelled resources for images are more plentiful (in fact, we are implicitly also using the even larger training set of the pre-trained ResNet). The availability of such large visual data sets allows us to train a strong external image tagger.

(B) *Image-speech pairs*: The Flickr8k Audio Captions Corpus [20], consisting of $\sim 8k$ images paired with 5 spoken captions each, amounting to a total of ~ 46 hours of speech data (~ 34 hours training, ~ 6 hours dev, and ~ 6 hours test speech data). The images in this set are disjoint from those in set A.

(C) *Speech-text pairs*: The spoken captions in the Flickr8k Audio Captions Corpus have written transcripts as well. We use subsets of these transcripts with varying sizes: from just ~ 21 minutes to the complete ~ 34 hours of labelled speech.

(D) *Human semantic relevance judgments*: For semantic speech retrieval evaluation, we use the human relevance judgments from [6]. 1000 utterances from the Flickr8k Audio Captions Corpus were manually annotated via Amazon Mechanical Turk with their semantic relevance for each of 67 query words. Each (utterance, keyword) pair was labeled by 5 annotators. We use both the majority vote of the annotators (as “hard labels”) and the actual number of votes (“soft labels”) for evaluation.

3.2. Implementation details

The *vision model* (Figure 1, left) is an ImageNet pre-trained ResNet-152 [21] (which is kept fixed), with a set of four 2048-unit fully connected layers added at the top, followed by a final softmax layer that produces posteriors for the N_{vis} tags. This image tagger is used to provide the visual supervision y_{vis} . The fully connected layers in the vision model are trained once on set A and kept fixed throughout the experiments.

The *speech model* (Figure 1, middle) consists of two branches, one with textual supervision and the other with visual supervision as targets. Except for the output layer, the architecture for these two networks is identical and exactly the same as the model in [6]. Each network consists of three convolutional layers, with the output max-pooled over time to get a 1024-dimensional embedding, followed by feedforward layers, with a final sigmoid layer producing either N_{vis} or N_{bow} scores in $[0, 1]$. The parameters for convolutional layers in these two branches are shared while the upper layer parameters are task-specific. Hereafter, the visually supervised branch and the textually supervised branch of this model are referred to as MTL-visSup and MTL-textSup respectively. The input speech is represented as MFCCs, zero-padded (or truncated²) to 8 seconds (800 frames).

For *representation learning* we need the dimensions for the learned speech and vision feature vectors to match. A two-layer ReLU feedforward network is used to transform the intermediate 2048-dimensional vision feature vectors to 1024 dimensions. The output of the shared branch of the speech model is used as the speech representation. Parameters of the additional layers are learned jointly with the speech model.

For MTL, each mini-batch is sampled from either set B or set C with probability proportional to the number of data points in each set, as in [22]. The size of set C varies, as described in 3.1. N_{vis} is kept fixed at 1000 and N_{bow} is one of 1k, 2k, 5k and 10k (each time keeping the most common content words in set A) and optimized along with other hyperparameters. Hyperparameters include α_{vis} , α_{bow} , m , n_{neg} , N_{bow} , batch size, and learning rate. Different settings are found to be optimal as the set C is varied.

²99.5% of the utterances are 8 s or shorter.

Adam optimization [23] is used for both branches with early stopping. For early stopping, we use F -score on the Flickr8k Audio Captions validation set with a detection threshold of 0.3. Since no development set is available for semantic retrieval, the F -score is evaluated for the task of exact retrieval (i.e., keyword spotting) using the 67 keywords from set D.

3.3. Baselines

We compare our proposed models to two baselines. The *visual baseline* has access to visual supervision alone. This baseline replicates the vision-speech model from [6], but is improved here by using a pre-trained ResNet-152 (rather than VGG-16) for fair comparison with our models. The visual baseline is equivalent to MTL-visSup, when no textual supervision is used. The *textual baseline* has access to just the textual supervision, and is equivalent to MTL-textSup when using no visual supervision. Different baseline scores are obtained as the set C changes.

3.4. Evaluation Metrics

At test time we have access to just the spoken utterances. For evaluation we use the output probability vectors of both, MTL-visSup and MTL-textSup. We evaluate semantic retrieval performance on set D (67 query words, 1000 spoken utterance search set). We measure performance using the following evaluation metrics, commonly used for retrieval tasks [24, 25]. *Precision at 10* ($P@10$) and *precision at N* ($P@N$) measure the precision of the top 10 and top N retrievals, respectively, where N is the number of ground-truth matches. *Average precision* is the area under the precision-recall curve as the detection threshold is varied. *Spearman’s rank correlation coefficient* (Spearman’s ρ) measures the correlation between the utterance ranking induced by the predicted probability vectors and the ground-truth “probability” vectors. The latter is approximated using the number of votes each query word gets for a given utterance. All of these metrics except for Spearman’s ρ are used for the hyperparameter tuning on exact keyword retrieval performance.

4. Results and Discussion

Figure 2 presents the semantic retrieval performance of the MTL model as the amount of text supervision is varied. For our MTL models, we can use either output (MTL-visSup or MTL-textSup) to perform the semantic retrieval task, and we explore the performance of both. We compare the MTL model results to those of the textual and visual baselines. Note that the visual baseline results are horizontal lines in each plot, as this baseline does not use any textual supervision. We present the results for Spearman’s ρ , which evaluates with respect to the “soft” human labels, and average precision, which uses the hard majority decisions. The trends for other hard label-based metrics ($P@10$, $P@N$) follow the same trend as average precision.

The first clear conclusion from these results is that visual grounding is still helpful even when some speech transcripts are available. This is, to our knowledge, the first time that this has been demonstrated for visually grounded models of speech semantics. In terms of average precision, the multitask model with both visual and textual supervision is only slightly better than the baseline text-supervised model when using larger amounts of transcribed speech (around 15 hours or more); when less text supervision is available, visual grounding becomes much more useful. At all levels of supervision, the best performance is obtained with one of the MTL models; depending on how much supervised data is available, either MTL-textSup or MTL-visSup

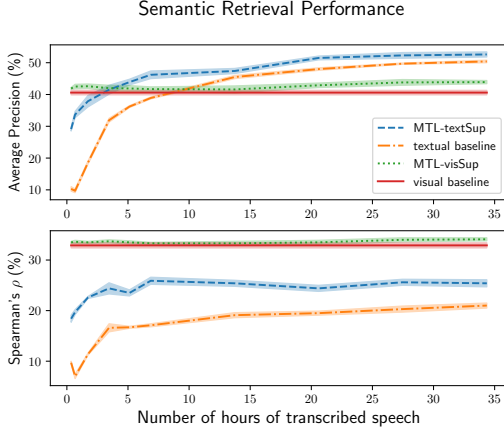


Figure 2: *Semantic retrieval performance, in terms of average precision and Spearman’s ρ . Shading around the curves indicates standard deviation over multiple runs with different random initializations.*

is preferred.

In terms of Spearman’s ρ , the trends are somewhat different from average precision. Here, the MTL-visSup outperforms MTL-textSup at all levels of supervision, and the visual baseline alone is almost as good as the MTL-visSup. This finding is in line with earlier work [6], which found that the visual baseline does particularly well in terms of Spearman’s ρ , which is more permissive of non-exact semantic matches. In some sense, Spearman’s ρ uses a more “complete” measure of the ground truth, since it considers the full range of human judgments rather than just the majority opinion. This is an interesting finding, but requires further study. For example, it could be the case that this domain is particularly well-suited to visual grounding because all of the speech is describing visual scenes. It will be interesting to study how these results change if we apply similar models to a wider range of speech domains.

Effect of representation loss. The results in Figure 2 use the supervised losses as well as the representation loss. Removing the latter reduces average precision by roughly 2-4% for most points on the curve, while for Spearman’s ρ it makes little difference. However, when tuning for exact retrieval on the development set, we find much larger gains from the representation loss, roughly 7-15% in average precision. This difference could be an artifact of tuning and testing on different tasks.

Effect of varying N_{bow} . We observe that a higher N_{bow} of 10k words is preferred at lower supervision, while at higher supervision having just 1k words in the output is best. Our interpretation is that having more words in the output than we will be evaluating on is equivalent to training the model on additional tasks which we do not care about at test time. This additional MTL-like setting has a regularization effect which is helpful in lower-resource settings, but as we have access to more and more labelled data, we no longer benefit from regularization.

Other modeling alternatives. In preliminary experiments, we also considered several other ways to combine the multiple sources of supervision: (1) by splitting the two branches at an earlier layer (2) by pre-training on one task and fine-tuning on another (target) task, or (3) by using a hierarchical multitask approach, inspired by prior related work [26–28], where the (more semantic) visual supervision is at a higher level and the

(exact) textual supervision is at a lower level. Our MTL model as presented in Figure 1 outperforms these other approaches.

Qualitative analysis. We visualize embeddings from the penultimate layer of the model by passing a set of isolated word segments through the model. Figure 3 shows 2-dimensional t-SNE [29] visualizations of these embeddings from both the textual baseline model and MTL-textSup near the low end of supervision. Qualitatively, the representation learned by the MTL model results in more distinct clusters than from the supervised baseline. We have also examined the retrieved utterances for these models and found that, at lower supervision, MTL-textSup outputs higher false positives than MTL-visSup due to acoustically similar words. For instance, for the query “tree”, MTL-textSup retrieves utterances containing the acoustically similar word “street”. This effect reduces as we train MTL-textSup on higher amount of text data.

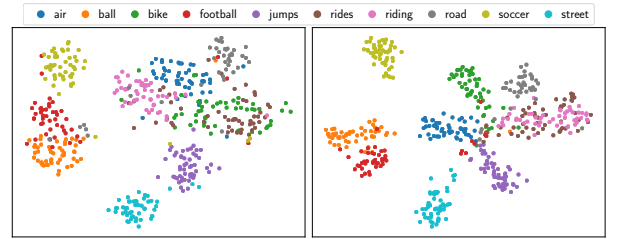


Figure 3: *T-SNE visualization of the representations learned with ~ 1.7 hours of transcribed speech, using the textual baseline (left) and MTL-textSup (right).*

5. Conclusion

Visual grounding has become a commonly used source of weak supervision for speech in the absence of textual supervision. Our motivation here was to, first, examine the contribution of visual grounding when some transcribed speech is also available and, second, explore how best to combine both visual and textual supervision for a semantic speech retrieval task. We proposed a multitask speech-to-keyword model that has both visually supervised and textually supervised branches, as well as an additional explicit speech-vision representation loss as a regularizer. We explored the performance of this model in a low-resource setting with various amounts of textual supervision, from none at all to almost 35 hours of transcribed speech. Our main finding is that visual grounding is indeed helpful even in the presence of some textual supervision. Experiments over a range of levels of supervision show that with different amounts of textual supervision, either the visual branch or textual branch of our model may perform better, but in all cases, joint training with both visual and textual supervision results in improved retrieval.

A limitation of the current work is that the set of queries and human judgments is small. A natural next step is to collect more human evaluation data and to consider a wider variety of queries, including multi-word queries. Another next step is to widen the types of speech domains that we consider, for example to explore whether visually grounded training can learn retrieval models that are applicable also to speech that is not describing a visual scene. On a technical level, there is room for more exploration of different types of multi-view representation losses (such as ones based on canonical correlation analysis [30, 31]), as well as more structured speech models that can localize the relevant words/phrases for a given query.

6. References

- [1] G. Synnaeve, M. Versteegh, and E. Dupoux, “Learning words from images and speech,” in *NIPS Workshop Learn. Semantics*, 2014.
- [2] D. Harwath, A. Torralba, and J. R. Glass, “Unsupervised learning of spoken language with visual context,” in *Proc. NIPS*, 2016.
- [3] G. Chrupała, L. Gelderloos, and A. Alishahi, “Representations of language in a model of visually grounded speech signal,” in *Proc. ACL*, 2017.
- [4] O. Scharenborg et al., “Linguistic unit discovery from multi-modal inputs in unwritten languages: Summary of the “Speaking Rosetta” JSALT 2017 Workshop,” in *Proc. ICASSP*, 2018.
- [5] H. Kamper and M. Roth, “Visually grounded cross-lingual keyword spotting in speech,” in *Proc. SLTU*, 2018.
- [6] H. Kamper, G. Shakhnarovich, and K. Livescu, “Semantic speech retrieval with a visually grounded model of untranscribed speech,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 27, no. 1, pp. 89–98, 2019.
- [7] L. Gelderloos and G. Chrupała, “From phonemes to images: Levels of representation in a recurrent neural model of visually-grounded language learning,” *Proc. COLING*, 2016.
- [8] D. Harwath and J. R. Glass, “Learning word-like units from joint audio-visual analysis,” in *Proc. ACL*, 2017.
- [9] D. Harwath, A. Recasens, D. Surís, G. Chuang, A. Torralba, and J. Glass, “Jointly discovering visual objects and spoken words from raw sensory input,” in *Proc. ECCV*, 2018.
- [10] D. Harwath and J. Glass, “Towards visually grounded sub-word speech unit discovery,” in *Proc. ICASSP*, 2019.
- [11] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” in *Proc. Interspeech*, 2017.
- [12] H. Kamper, A. Anastassiou, and K. Livescu, “Semantic query-by-example speech search using visual grounding,” in *Proc. ICASSP*, 2019.
- [13] J. S. Garofolo, C. G. Auzanne, and E. M. Voorhees, “The TREC spoken document retrieval track: A success story,” in *Proc. CBMIA*, 2000.
- [14] L.-s. Lee, J. R. Glass, H.-y. Lee, and C.-a. Chan, “Spoken content retrieval—beyond cascading speech recognition with text retrieval,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 23, no. 9, pp. 1389–1420, 2015.
- [15] C. Chelba, T. J. Hazen, and M. Saraclar, “Retrieval and browsing of spoken content,” *IEEE Signal Proc. Mag.*, vol. 25, no. 3, 2008.
- [16] K. M. Hermann and P. Blunsom, “Multilingual distributed representations without word alignment,” in *Proc. ICLR*, 2014.
- [17] W. He, W. Wang, and K. Livescu, “Multi-view recurrent neural acoustic word embeddings,” in *ICLR*, 2017.
- [18] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft COCO: Common objects in context,” in *Proc. ECCV*, 2014.
- [19] P. Young, A. Lai, M. Hodosh, and J. Hockenmaier, “From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions,” *Trans. ACL*, vol. 2, pp. 67–78, 2014.
- [20] D. Harwath and J. R. Glass, “Deep multimodal semantic embeddings for speech and images,” in *Proc. ASRU*, 2015.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016.
- [22] V. Sanh, T. Wolf, and S. Ruder, “A hierarchical multi-task approach for learning embeddings from semantic tasks,” *arXiv:1811.06031*, 2018.
- [23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.
- [24] T. J. Hazen, W. Shen, and C. White, “Query-by-example spoken term detection using phonetic posteriorgram templates,” in *Proc. ASRU*, 2009.
- [25] Y. Zhang and J. R. Glass, “Unsupervised spoken keyword spotting via segmental DTW on Gaussian posteriorgrams,” in *Proc. ASRU*, 2009.
- [26] A. Søgaard and Y. Goldberg, “Deep multi-task learning with low level tasks supervised at lower layers,” in *Proc. ACL*, 2016.
- [27] S. Toshniwal, H. Tang, L. Lu, and K. Livescu, “Multitask Learning with Low-Level Auxiliary Tasks for Encoder-Decoder Based Speech Recognition,” in *Proc. Interspeech*, 2017.
- [28] R. Sanabria and F. Metze, “Hierarchical Multi Task Learning With CTC,” in *Proc. SLT*, 2018.
- [29] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [30] W. Wang, R. Arora, K. Livescu, and J. Bilmes, “On deep multi-view representation learning,” in *Proc. ICML*, 2015.
- [31] Y. Gong, Q. Ke, M. Isard, and S. Lazebnik, “A multi-view embedding space for modeling internet images, tags, and their semantics,” *International journal of computer vision*, vol. 106, no. 2, pp. 210–233, 2014.