

Semantic speech retrieval with a visually grounded model of untranscribed speech

Herman Kamper, Gregory Shakhnarovich, and Karen Livescu

Abstract—There is growing interest in models that can learn from unlabelled speech paired with visual context. This setting is relevant for low-resource speech processing, robotics, and human language acquisition research. Here we study how a visually grounded speech model, trained on images of scenes paired with spoken captions, captures aspects of semantics. We use an external image tagger to generate soft text labels from images, which serve as targets for a neural model that maps untranscribed speech to (semantic) keyword labels. We introduce a newly collected data set of human semantic relevance judgements and an associated task, *semantic speech retrieval*, where the goal is to search for spoken utterances that are semantically relevant to a given text query. Without seeing any text, the model trained on parallel speech and images achieves a precision of almost 60% on its top ten semantic retrievals. Compared to a supervised model trained on transcriptions, our model matches human judgements better by some measures, especially in retrieving non-verbatim semantic matches. We perform an extensive analysis of the model and its resulting representations.¹

Index Terms—Visual grounding, multimodal modelling, speech retrieval, semantic retrieval, keyword spotting.

I. INTRODUCTION

THE last few years have seen great advances in automatic speech recognition (ASR). However, current methods require large amounts of transcribed data, which are available only for a small fraction of the world’s languages [1]. This has prompted work on speech models that, instead of using exact transcriptions, can learn from weaker forms of supervision, e.g., known word pairs [2], [3], translation text [4]–[6], or unordered word labels [7]. The motivation for much of this work is that, even when high-quality ASR is infeasible, it may still be possible to learn low-resource models for practical tasks like retrieval and keyword prediction.

In this work we consider the setting where untranscribed speech is paired with images. Such visual context could be used to ground speech when it is not possible to obtain transcriptions, e.g., for endangered or unwritten languages [8]. In robotics, co-occurring audio and visual signals could be combined to learn new commands [9]–[11].

Our work builds on a line of recent studies [12]–[19] that use natural images of scenes paired with spoken descriptions.

H. Kamper is with the Department of E&E Engineering, Stellenbosch University, South Africa (email: kamperh@sun.ac.za).

G. Shakhnarovich and K. Livescu are with TTI-Chicago, USA (email: greg@ttic.edu, klivescu@ttic.edu)

¹The collected data set of semantic labels is available at: https://github.com/kamperh/semantic_flickr_audio.

The code recipe for the neural networks developed in this work is available at: https://github.com/kamperh/recipe_semantic_flickr_audio.

Neither the spoken nor visual input is labelled. Most approaches map the images and speech into some common space (§II-C), allowing images to be retrieved using speech and vice versa. Although useful, such models cannot predict (written) labels for the input speech.

Here we present an expansion of our work in [20], where we proposed a model that can make such text label predictions. A trained visual tagger is used to obtain soft text labels for each training image, and these are used as targets for a neural network that maps speech to keyword labels. The resulting model can be used to predict which (written) words are present in an unseen spoken utterance. Without observing any transcriptions, the model can be used as a keyword spotter, predicting which utterances in a search collection contain a given written keyword. In an initial qualitative analysis, we observed that the model often confuses semantically related words such as ‘man’ and ‘person’; these count as errors in keyword spotting, but should be useful in *semantic* search applications.

Our primary aim here is to expand our work in [20] by performing an extensive analysis to see what aspects of semantics are captured by the model. To do so formally, we use the task of *semantic speech retrieval*, where the aim is to retrieve all utterances in a speech collection that are semantically relevant to a given query keyword, irrespective of whether that keyword occurs exactly in an utterance or not. E.g., given the query ‘children’, the goal is to return not only utterances containing the word ‘children’, but also utterances about children, like ‘young boys playing soccer in the park’. There has been some work on this and related tasks (§II-E), but typically in higher-resource settings and none using visual supervision. Note that visual supervision could actually prove more useful than transcriptions for semantic retrieval in some cases (and we make findings in this direction). To our knowledge, despite the prior work, speech data with soft semantic relevance judgements does not exist. This limits analyses of semantics, which are often ambiguous. We therefore collect and release a new data set with semantic labels from multiple annotators. Using this data, we present an extensive analysis of an updated version of the model of [20], and compare it to several new alternative models for the task of semantic speech retrieval.

Our main contributions are to extensively analyse the visually grounded speech modeling approach in the context of a formal semantic task: We provide the new data set, compare model performance to human judgements, study the learned

representations, and analyse the performance of the model against multiple alternatives. We find that the predictions from the model match well with human judgements, leading to competitive retrieval scores, particularly in retrieving non-exact semantic matches. Specifically, the model even outperforms a supervised model trained on transcriptions when searching for utterances that are semantic but not verbatim matches to the keyword. We conclude that visual context can play an important role in capturing semantics in speech, especially but not only in the absence of textual supervision.

II. RELATED WORK

A. Language acquisition

Cognitive scientists have long been interested in how infants use sensory input to learn the mapping of words to real-world entities [21]–[23], with computational models providing one way to specify and test theories. Roy and Pentland [24] and Yu and Ballard [25] were some of the first to consider real speech input, followed by more recent work such as [26]–[28]. However, these studies simplify the problem by using discrete labels to represent the visual context [29], the spoken input [30], or both [31], [32]; infants do not have access to such idealised input. Our model operates on natural images paired with real unlabelled speech. One assumption we make is that a visual tagger is available for processing training images. Our focus is not on cognitive modelling, but this assumption is linked to the question of whether visual category acquisition precedes word learning in infants [33]. We leave the cognitive implications of our model for future work.

Also motivated by the problem of modelling language acquisition, there is growing interest in the speech community in models that can learn directly from unlabelled speech [34], [35]. These studies have the additional motivation that such systems could be applied in settings where it is difficult or impossible to obtain transcriptions. The goal here is to learn exclusively from audio, i.e. without any supervision signal, even from another modality. Tasks include acoustic unit discovery [36], [37], query-by-example search [38], unsupervised term discovery [39], and unsupervised segmentation and clustering [40], [41]. See [42] for a complete literature review. Although our work is related to these studies, we focus on the setting where visual context is available for grounding.

B. Joint modelling of images and text

Joint modelling of images and text has received much recent attention. One approach is to map images and text into a common vector space where related instances are close to each other, e.g., for text-based image retrieval [43]–[46]. Image captioning has also been studied extensively, where the goal is to produce a natural language description for a visual scene [47]–[50]. Most recent approaches use a convolutional neural network (CNN) to convert an input image to a latent representation, which is then fed to a recurrent neural network to produce a sequence of output words [51]–[56]. Our work uses spoken rather than written language, and we consider a semantic speech retrieval task.

There has also been work on using vision to explicitly capture aspects of text semantics. Semantics are difficult to annotate, so most studies evaluate models using soft human ratings on tasks such as word similarity [57]–[59], word association [60], [61] or concept categorisation [62]. We also use human responses, but for the task of semantic speech retrieval.

C. Joint modelling of vision and speech

Recent work has shown that ASR can be improved by including additional visual features from the scene in which speech occurs [63], [64]. These studies consider fully supervised ASR, and typically the scene is not described by the speech but is complementary to it. Our aim instead is to use the visual modality to learn from matching but *untranscribed* speech.

This is the setting considered by Synnaeve et al. [12] and Harwath et al. [13], who used natural images of scenes paired with unlabelled spoken descriptions to learn neural mappings of images and speech into a common space. This approach allows images to be retrieved using speech and vice versa. This is useful, e.g., in applications for tagging mobile phone images with spoken descriptions [65], [66]. The joint neural mapping approach has subsequently been used for spoken word and object segmentation [67], and the learned representations have been analysed in a variety of ways [8], [15]. Despite these developments, this prior work does not give an explicit mapping of speech to textual labels.

The model we proposed in [20] can make such labelled predictions: a visual tagger is used to obtain soft labels for a given training image, and a neural network is then trained to map input speech to these targets. The resulting speech model attempts to predict which (written) words are present in a spoken utterance, ignoring the order and quantity of words. The model can be used as a keyword spotter, retrieving utterances in a speech collection that contain a given keyword. It was found that the model also retrieves semantic matches for a query, not only exact matches. However, we did not formalise the semantic search task, compare results to human judgements, or extensively analyse the performance of the model against multiple other systems. Here we present an extensive analysis using several alternative models for semantic speech retrieval.

D. Semantic text retrieval

In textual information retrieval, the task is to find text documents in a collection that are relevant to a given written keyword, irrespective of whether the keyword occurs exactly in the document. This can be accomplished using *query expansion*, where additional words or phrases similar in meaning to the query are used to match relevant documents [68], [69]. Classic approaches use co-occurrences of words, or resources such as WordNet [70], to expand the query list, while recent methods use word embeddings [71], [72]. We also consider a semantic search task, but on speech rather than text. We do use text-based methods to obtain an upper bound on performance using the transcriptions of the speech collection (§V-B).

E. Keyword spotting and (semantic) speech search

In the speech research literature, keyword spotting and spoken term detection involve retrieving utterances in a search collection that contain exact spoken instances of a given written keyword [73]–[75]. Here we specifically consider *semantic* rather than exact retrieval. There has been some work on semantic speech retrieval [76]–[79]; the most common approach is to cascade ASR with text-based retrieval. This, again, requires large amounts of labelled data with which to train an accurate ASR system. There have been some attempts to do semantic retrieval in the absence of transcriptions by extending unsupervised speech modelling approaches (described in §II-A). In [80], an unsupervised term detection system was used to automatically discover reoccurring words or phrases. Given a spoken query, the system then expanded the query using other commonly co-occurring discovered terms, thereby allowing retrieval of semantically related words (other than the one spoken). In [81], a topic model was applied to discovered words in a similar way. Here we are also doing semantic speech retrieval in the absence of transcriptions. Our task is different, however, in that our queries are textual keywords; this is not possible using a model trained without some grounding signal (which is why the above studies performed query-by-example search). We are considering the case where we have limited supervision in the form of visual context. This is a first, to our knowledge.

III. THE VISUALLY GROUNDED SPEECH MODEL

Given a corpus of parallel images and spoken captions, neither with textual labels, we train a spoken keyword prediction model using a visual tagging system to produce soft labels for the speech network. Originally, we used the model for keyword spotting [20], with the output of the model used to predict whether a given written keyword occurs exactly in a test utterance. Here, we also interpret the output of the model as a score for a keyword, but instead of looking for an exact match, we use it for predicting semantic relevance.

The overall approach is illustrated in Figure 1. Training image I is paired with a spoken caption $X = x_1, x_2, \dots, x_T$, where each frame x_t is an acoustic feature vector, e.g. Mel-frequency cepstral coefficients (MFCCs) [82]. We use an external vision system to tag I with soft textual labels, giving targets \hat{y}_{vis} to train the speech network $f(X)$. The resulting network $f(X)$ can then predict which keywords are present for a given utterance X , disregarding the order, quantity, or locations of the keywords in the input speech. The possible keywords (the vocabulary) are implicitly specified by the visual tagger, and no transcriptions are used during training. When applying the trained $f(X)$ for keyword spotting or semantic retrieval, test speech alone is used without any visual input.

A. Detailed model description

In typical speech applications like keyword spotting, supervision consists of text labels of the keywords that appear in the speech: for training utterance X , we could construct a multi-hot vector $\mathbf{y}_{\text{bow}} \in \{0, 1\}^W$, with W the vocabulary size and each dimension $y_{\text{bow},w}$ a binary indicator for whether word w occurs

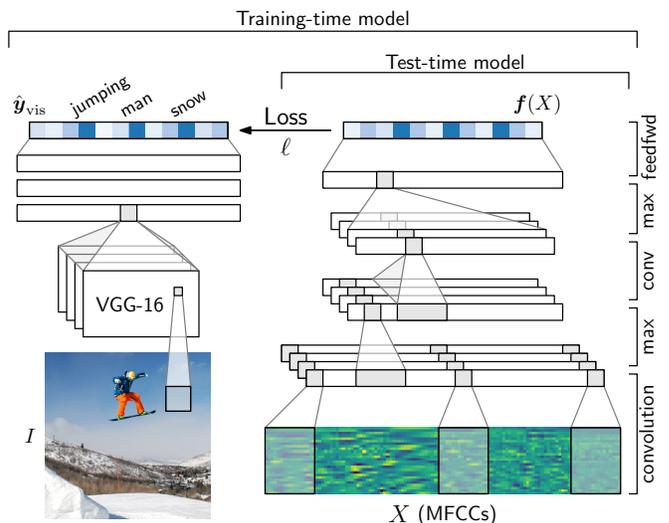


Fig. 1. For training, an external visual tagger produces soft tags for image I , which serve as targets for the speech network fed with spoken caption X . In testing, the speech network is given unseen speech (no image) and the output $f(X)$ is used for semantic retrieval of a textual keyword.

in X . The order and quantity of words are ignored in this type of bag-of-words (BoW) labelling [7]. Instead of a transcription for X , we only have the paired image I . We use a multi-label visual tagging system which, instead of binary indicators, produces soft targets $\hat{y}_{\text{vis}} \in [0, 1]^W$, with $\hat{y}_{\text{vis},w} = P_{\gamma}(w|I)$ the estimated probability of word w being present in image I under vision model parameters γ . In Figure 1, \hat{y}_{vis} would ideally be close to 1 for w corresponding to words such as ‘jumping’, ‘man’, ‘snow’ and ‘snowboard’, and close to 0 for irrelevant dimensions. For semantic retrieval, we therefore interpret \hat{y}_{vis} as a bag-of-visual (semantic) tags. This visual tagger is fixed: during training of the speech network $f(X)$, as described below, vision parameters γ are never updated.

Given \hat{y}_{vis} as target, we train the speech model $f(X)$ (Figure 1, right). This model (parameters θ) consists of a convolutional neural network (CNN) over the speech X with a final sigmoidal layer so that $f(X) \in [0, 1]^W$. We interpret each dimension of the output as $f_w(X) = P_{\theta}(w|X)$. Note that $f(X)$ is not a distribution over the output vocabulary, since any number of keywords can be present in an utterance; rather, it is a multi-label classifier where each dimension $f_w(X)$ can have any value in $[0, 1]$. We train the speech model using the summed cross-entropy loss, which (for a single training example) is:

$$\ell(f(X), \hat{y}_{\text{vis}}) = - \sum_{w=1}^W \{ \hat{y}_{\text{vis},w} \log f_w(X) + (1 - \hat{y}_{\text{vis},w}) \log [1 - f_w(X)] \} \quad (1)$$

If we had $\hat{y}_{\text{vis},w} \in \{0, 1\}$, as in \mathbf{y}_{bow} , this would be the summed log loss of W binary classifiers. A pre-trained visual network was also used to provide supervision for another modality using a similar loss in [83], where video was paired with general audio (not speech).

B. The visual tagging system

In image classification, the task is to choose one (object) class from a closed set [84], while in image captioning, the goal is to produce a natural language description of a scene (§II-B). In contrast to both these tasks, we require a visual tagging system [85]–[87] that predicts an unordered set of words (nouns, adjectives, verbs) that accurately describe aspects of the scene (Figure 1, left). This is a multi-label binary classification task.

We train our visual tagger on data from the Flickr30k [88] and MSCOCO [89] data sets, which consist of images each with five written captions. We combine Flickr30k with the training portion of MSCOCO, and remove any images that occur in the parallel image-speech data used in our experiments (§V). The result is a training set of around 106k images, significantly more than the 25k used in [20]. For each image, a single target BoW vector \mathbf{y}_{vis} is constructed by combining all five captions after removing stop words. Element $y_{\text{vis},w}$ is an indicator for whether word w occurs in any of the five captions for that image, where w is one of the $W = 1000$ most common content words in the combined set of image captions.

We follow the common practice of using a pre-trained visual representation for processing the images (Figure 1, left). Specifically, we use VGG-16 [90], trained on around 1.3M images [84]; we replace the final classification layer with four 2048-unit rectified linear unit (ReLU) layers, followed by a final sigmoidal layer for predicting word occurrence. The visual tagger, with parameters γ , is then trained on the combined Flickr30k and MSCOCO data using the summed log loss (1) with arguments $\ell(\hat{\mathbf{y}}_{\text{vis}}, \mathbf{y}_{\text{vis}})$. The VGG-16 parameters are fixed; only the additional fully connected layers are updated.

Previous vision-speech models [12]–[14] also employ pre-trained visual representations. Here we take this approach even further by using the textual classification output of a trained vision system. Although we train (and then fix) the visual tagger ourselves, we ensure that none of the training data overlaps with the parallel image-speech data used in our speech model, so the model does not obtain even indirect access to text labels.

IV. SEMANTIC SPEECH RETRIEVAL

We are interested in whether our model can be used to determine the semantic concepts present in a speech utterance. I.e., can we use the model to search a speech collection for mentions of a particular semantic concept? We formalise this task and collect a new data set for evaluation.

A. Task description

We consider the task of *semantic speech retrieval*. Instead of matching keywords exactly, as is typical in keyword spotting (§II-E), the aim is to retrieve all utterances that are semantically relevant, irrespective of whether the keyword occurs in the utterance or not. E.g., for the query ‘sidewalk’, a model should return not only utterances containing the word exactly, but also speech like ‘an old couple window-shopping on a Paris street.’ As noted, there has been some work on semantic speech retrieval, but typically in higher-resource settings and none using visual supervision (§II-E). Furthermore, the new data set described below contains labels from multiple annotators, allowing for a fine-grained analysis.

TABLE I
THE PROPORTION OF (SENTENCE, KEYWORD) PAIRINGS SELECTED BY A GIVEN NUMBER OF ANNOTATORS.

Count	0	1	2	3	4	5
Proportion (%)	83.3	6.9	3.0	2.6	2.5	1.6

B. Data set

We extend the corpus of [91], which consists of parallel images and spoken captions. The data is transcribed, but not semantically labelled. For a subset of the speech in the corpus, we use Amazon Mechanical Turk (AMT) to collect semantic labels from human annotators.

As our keywords, we select a set of 70 random words from the transcriptions of the training portion of the corpus, ignoring stop words. The test portion of the spoken caption data consists of 1000 images, each with five spoken captions; we collect semantic labels for one randomly selected spoken caption from each of these 1000 images. A single AMT job consists of the transcription of a single utterance (describing a scene) with a list of seven potential keywords from which an annotator could select any number, as illustrated in Figure 2. To cover all 70 keywords, a given sentence is repeated over ten jobs. Since the question of semantic relatedness (between a given sentence and keyword) is inherently ambiguous, we have five workers annotate each utterance. Annotators were not shown any images but only the utterance transcriptions and candidate keywords, so the task is the same as in other work on semantic speech retrieval (§II-E).

To analyse annotator agreement, we consider the number of annotators (between zero and five) that selected a given keyword for each sentence. Table I gives the proportion of (sentence, keyword) pairs selected by a given number of annotators. Annotators agree most often (83.3% of the time) about the absence of a keyword for a given sentence, while there are very few examples that all annotators agree are semantically matched (1.6%).

To evaluate a semantic keyword retrieval model against the human annotations, one option is to combine the human judgements into a single hard label. On the other hand, the fact that there is a wide range of opinions among the human annotators indicates that semantic relevance may be inherently “soft”, motivating evaluation by comparing against the proportion of annotators that agree with a given label. In work using vision to explicitly aspects of semantics in text, soft human ratings were also used (§II-B). In our experiments we consider both hard and soft options.

For the sentence below, select all words that could be used to search for the described scene.

Sentence: a skateboarder in a light green shirt.

- dogs beach children white swimming
 wearing skateboard None of the above

Fig. 2. An example AMT job for semantically annotating a transcription of a spoken sentence with keywords. In this case, ‘skateboard’ was selected by all five annotators and ‘wearing’ by four.

To obtain a hard label of whether a sentence is semantically relevant or irrelevant given a keyword, we take the majority decision: if three or more annotators selected that keyword, we label that sentence as relevant for that keyword, otherwise we label it as irrelevant. Under this assignment, 95.8% of annotators agree with the hard decision (bearing in mind the skew towards negative assignments). We removed three keywords with very poor agreement, leaving 67 keywords.

The result of our data collection effort is a data set of 1000 spoken utterances, each annotated with a set of keywords that could be used to search for that utterance.

V. EXPERIMENTAL SETUP AND EVALUATION

We train our model on the corpus of parallel images and spoken captions of [91], containing 8000 images with five spoken captions each. The audio comprises around 37 hours of active speech. The data comes with train, development and test splits containing 30 000, 5000 and 5000 utterances, respectively. As described above, we obtained semantic labels for 1000 of the test utterances. We parameterise the speech audio as 13 MFCCs with first and second order derivatives, giving 39-dimensional input vectors.² 99.5% of the utterances are 8 s or shorter; utterances longer than this are truncated while utterances shorter than this are zero-padded to 8 s (800 frames). We deal with the variable duration of utterances by pooling all units at the output of convolutional layers over time (see below).³

Training images are passed through the visual tagger (§III-B), producing soft targets \hat{y}_{vis} for training the keyword prediction model $f(X)$ on the unlabelled speech (§III-A), as shown in Figure 1. We refer to the resulting model as VISIONSPEECHCNN, which during testing is presented only with spoken input. We are interested in *semantic* speech retrieval (§IV), but labels for this task are only available for test utterances. We therefore optimise the hyper-parameters of the model using *exact* keyword spotting on development data using the same set of 67 keywords; each spoken caption in the development data is treated as a separate test item. VISIONSPEECHCNN has the following resulting structure: 1-D ReLU convolution with 64 filters over 9 frames; max pooling over 3 units; 1-D ReLU convolution with 256 filters over 10 units; max pooling over 3 units; 1-D ReLU convolution with 1024 filters over 11 units; max pooling over all units; 3000-unit fully-connected ReLU; and the 1000-unit sigmoid output. A stride of one unit is used for all convolution operations and there is no overlap between windows in the max pooling operations. Based on experiments on development data, we train for a maximum of 25 epochs with early stopping using Adam optimization [92] with a learning rate of $1 \cdot 10^{-4}$ and a batch size of eight.⁴

²We also tried filterbanks; MFCCs always gave similar or slightly better performance on development data.

³We also experimented with using the actual durations of the utterances, but found minimal difference. Using the actual durations incurs an additional computational cost, whereas our approach can ignore the zero-padded regions through the global max pooling operation over time.

⁴The code recipe for the neural networks developed in this work is available at: https://github.com/kamperh/recipe_semantic_flickraudio.

A. Evaluation

To use VISIONSPEECHCNN for semantic retrieval, we use its output $f_w(X) \in [0, 1]$ as a score for how relevant an utterance X is given the keyword w . The baseline and cheating models (described below) similarly predict a relevance score for each utterance given a specific keyword.

We compare a model’s predictions to semantic labels obtained from human annotators (§IV-B) using several metrics. To obtain a hard labelling from a model, we set a threshold α , and label all keywords for which $f_w(X) > \alpha$ as relevant. By comparing this to the ground truth semantic labels (according to majority annotator agreement), precision and recall can be calculated; to measure performance independent of α , we report **average precision (AP)**, the area under the precision-recall curve as α is varied. Instead of using the hard ground truth semantic labels, the soft scores $f_w(X)$ can also be compared directly to the number of annotators that selected the keyword w for utterance X : we use **Spearman’s ρ** to measure the correlation between the rankings of these two variables, as is common in work on word similarity [93], [94]. The remaining metrics are standard in (exact) keyword spotting, based on how a model ranks utterances in the test data from most to least relevant for each keyword [95], [96]: **precision at ten ($P@10$)** is the average precision of the ten highest-scoring proposals; **precision at N ($P@N$)** is the average precision of the top N proposals, with N the number of true occurrences of the keyword; and **equal error rate (EER)** is the average error rate at which false acceptance and rejection rates are equal.

Apart from Spearman’s ρ , all these metrics can also be used to evaluate exact keyword spotting.

B. Baselines and cheating models

We consider a number of baselines as well as “cheating” models which use idealised information not available to VISIONSPEECHCNN.

Prior-based baselines: TEXTPRIOR uses the unigram probability of each keyword estimated from the transcriptions of the training portion of the spoken captions corpus. This will indicate how much better our model does than simply hypothesising common words. Similarly, VISIONTAGPRIOR is obtained by passing all training images through the trained visual tagger (§III-B), and then taking the average over all images. It therefore always predicts common visual labels.

VISIONCNN: VISIONSPEECHCNN is trained to predict soft visual tags. One question is whether it therefore learns to ignore any aspect of the acoustics that does not contribute to predicting the visual target. The model VISIONCNN is an attempt to test this: as the representation for each test utterance, it passes through the visual tagger the *true* image paired with that utterance. Since it uses ideal information, it is our first cheating model. If VISIONSPEECHCNN were to perfectly predict image tags (ignoring acoustics that do not contribute to visual prediction), then VISIONCNN would be an upper bound on performance. But in reality our model could do better or worse than VISIONCNN, since the speech contains some information not in the images and training does not generalise perfectly.

TABLE II

THE RETRIEVED UTTERANCE RATED HIGHEST BY VISIONSPEECHCNN FOR A SELECTION OF KEYWORDS. THE NUMBER OF ANNOTATORS (OUT OF FIVE) THAT SELECTED THE KEYWORD FOR THAT UTTERANCE IS SHOWN, WITH * INDICATING AN INCORRECT SEMANTIC RETRIEVAL ACCORDING TO THE MAJORITY LABELLING.

Keyword	Top retrieved utterance	Count
jumps	biker jumps off of ramp	5 / 5
ocean	man falling off a blue surfboard in the ocean	5 / 5
race	a red and white race car racing on a dirt racetrack	5 / 5
snowy	a skier catches air over the snow	5 / 5
bike	a dirt biker rides through some trees	4 / 5
children	a group of young boys playing soccer	4 / 5
riding	a surfer rides the waves	4 / 5
young	a little girl in a swing laughs	4 / 5
field	two white dogs running in the grass together	3 / 5
swimming	a woman holding a young boy slide down a water slide into a pool	3 / 5
carrying	small dog running in the grass with a toy in its mouth	2 / 5 *
face	a man in a gray shirt climbs a large rock wall	2 / 5 *
large	a group of people on a zig path through the mountains	1 / 5 *
sitting	a baby eats and has food on his face	1 / 5 *
hair	two women and a man smile for the camera	0 / 5 *

SUPERVISED BOWCNN: Instead of soft targets from an image tagger, the SUPERVISED BOWCNN cheating model uses transcriptions to obtain hard BoW supervision (§III-A): \mathbf{y}_{bow} targets are constructed for the 1000 most common words in the transcriptions of the 30 000 speech training utterances (ignoring stop words) and the loss (1) is used for training, $\ell(\mathcal{f}_{\text{bow}}(X), \mathbf{y}_{\text{bow}})$. Other than ideal supervision, the model has the same structure and training procedure as VISIONSPEECHCNN. This model allows us to explore how the visual supervision differs from typical textual transcriptions for semantic concept learning.

Text-based cheating models: Suppose we had a perfect ASR system, converting input speech to text without errors. How well could we do at semantic speech retrieval using this error-free text? To answer how this cascaded approach would do in an ideal setting, we consider two text-based semantic retrieval methods (§II-D) using transcriptions of the speech. The first is based on WuP similarity, named after Wu and Palmer [97], which scores the semantic relatedness between two words according to the path length between them in the WordNet lexical hierarchy [70]. For our task, the TEXTWUP model is based on the closest match (in WuP) between a keyword and each of the words in a transcribed utterance. Our second text-based cheating method is based on word embeddings. Specifically, we use the PARAGRAMXXL embedding method of [98], [99], developed for semantic sentence similarity prediction. For our task, the TEXTPARAGRAM cheating model calculates the cosine similarity between a keyword embedding and the PARAGRAM sentence embedding of an utterance.

VI. EXPERIMENTAL RESULTS AND ANALYSIS

For a qualitative view, Table II shows the top retrievals when we use VISIONSPEECHCNN to do semantic speech retrieval for a selection of keywords. The number of annotators that marked

the utterance as relevant is also shown; * indicates incorrect retrievals according to the ground truth (i.e. majority) semantic labelling. Out of the 15 results shown, ten retrievals are correct.

The quantitative metrics for exact and semantic keyword retrieval for VISIONSPEECHCNN and all the baseline and cheating models are shown in Table III. In both exact and semantic retrieval, VISIONSPEECHCNN outperforms the baseline models across all metrics. The baseline models, VISIONSPEECHCNN, and VISIONCNN all perform better at semantic than at exact retrieval. In contrast, the transcription-based cheating models (rows 5 to 7) perform better on $P@10$, but worse on all other semantic search metrics. $P@10$ only measures precision of the highest ranked utterances, while the other metrics combine precision and recall; thus, the transcription-based cheating models struggle to retrieve semantic matches compared to exact matches, while VISIONSPEECHCNN and VISIONCNN recall more semantic matches. In terms of absolute performance, the transcription-based models still perform better at semantic speech retrieval on the metrics based on hard ground truth labels. However, for Spearman’s ρ , which gives credit even if a prediction does not match the majority of annotations, VISIONCNN outperforms all other models, followed closely by VISIONSPEECHCNN. Visual context is clearly beneficial in matching soft human ratings. Next, we further analyse the models by addressing the following questions.

Does VISIONSPEECHCNN only output common words?

The baseline models (rows 1 and 2) respectively assign scores to keywords according to their occurrence in the training transcriptions and the average visual tagger output. VISIONSPEECHCNN outperforms both across all metrics.

Does the model do more than map acoustics to images? One possibility is that VISIONSPEECHCNN might learn to in effect ignore signals in the acoustics that are irrelevant to producing the visual output. To see if this is so, we compare with VISIONCNN (row 4), which represents a test utterance by passing the paired test image through the visual tagger. VISIONSPEECHCNN outperforms VISIONCNN across all metrics, except for Spearman’s ρ . This indicates that VISIONSPEECHCNN does not simply reproduce the output from the visual tagger (used to supervise it); it actually achieves superior performance for all exact keyword spotting metrics and most of the semantic metrics. Again, Spearman’s ρ takes the actual annotator counts into account, and here VISIONCNN performs best.

Is there a benefit to using visual context over transcriptions? This work is motivated by settings in which text transcriptions are not available while images may be. However, it is interesting to consider whether there is some benefit to the training images beyond serving as weak labels. Might the visual grounding actually provide better supervision than text for some purposes?

We compare rows 3 and 5 of Table III. In keyword spotting, VISIONSPEECHCNN lags behind SUPERVISED BOWCNN (trained on ground truth transcriptions). However, when moving to semantic speech retrieval, VISIONSPEECHCNN improves on almost all metrics, while SUPERVISED BOWCNN performs worse on most scores relative to exact keyword spotting. VISIONSPEECHCNN also outperforms SUPERVISED BOWCNN on Spearman’s ρ , so it matches better the soft human judgements.

TABLE III

KEYWORD SPOTTING AND SEMANTIC SPEECH RETRIEVAL PERFORMANCE FOR VISIONSPEECHCNN (ROW 3), COMPARED AGAINST THE BASELINE (ROWS 1 AND 2) AND CHEATING (ROWS 4 TO 7) MODELS. BOLDFACE INDICATES BOTH THE TOP-SCORING NON-CHEATING MODEL (ROWS 1 TO 3) AS WELL AS THE BEST CHEATING MODEL (ROWS 4 TO 7) FOR EACH OF THE METRICS.

Model	Exact keyword spotting (%)				Semantic speech retrieval (%)					
	$P@10$	$P@N$	EER	AP	$P@10$	$P@N$	EER	AP	Spear. ρ	
Baseline models:										
1. TEXTPRIOR	2.8	3.4	50.0	8.7	6.1	7.0	50.0	11.4	10.8	
2. VISIONTAGPRIOR	2.8	3.4	50.0	7.0	6.1	7.0	50.0	13.6	12.5	
3. VISIONSPEECHCNN	38.5	30.8	19.6	26.9	58.8	39.7	23.9	39.4	32.4	
Cheating models:										
4. VISIONCNN	31.0	26.2	22.1	22.2	54.2	38.9	22.8	37.4	33.8	
5. SUPERVISEDBOWCNN	84.9	74.7	5.6	87.3	88.1	50.3	23.8	51.3	21.9	
6. TEXTWUP	65.4	67.3	2.6	75.2	80.3	63.0	19.4	60.9	25.2	
7. TEXTPARAGRAM	80.0	72.1	3.5	67.7	88.8	64.0	14.3	60.1	31.6	

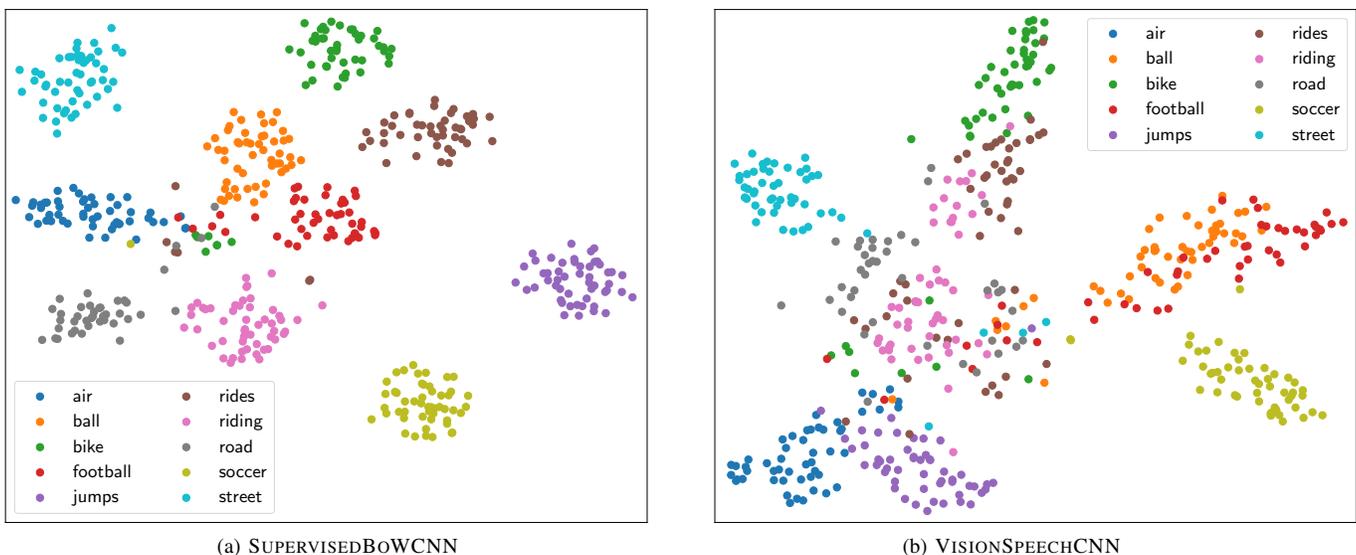


Fig. 3. t-SNE visualisations of acoustic embeddings of isolated words for ten keyword types. The embeddings were obtained from penultimate bottleneck layers for words from the development data.

This indicates that visual grounding actually provides better supervision than text for some purposes.

Despite the benefit of visual supervision, SUPERVISEDBOWCNN still performs better on semantic speech retrieval measured against hard labels (in absolute scores). We now show that this is because many exact matches are also semantic matches, which SUPERVISEDBOWCNN predicts well, while VISIONSPEECHCNN actually performs better on non-verbatim semantic matches. To do this, we consider separately the contributions of exact and exclusively semantic matches for the semantic $P@N$ metric (which actually also captures recall, since N is the number of true occurrences of a keyword). Standard $P@N$ cannot be broken into separate components, since it is averaged over keywords. We therefore define an unweighted $P@N^* = \frac{\sum_w c_w}{\sum_w N_w}$, where N_w is the number of occurrences of keyword w in the ground truth labels, and c_w is the number of correct utterances in the top N_w predictions. Count c_w can then be broken into exact and semantic matches, such that $P@N^* = P@N^*_{\text{exact}} + P@N^*_{\text{sem}}$ (here $P@N^*_{\text{sem}}$

TABLE IV
UNWEIGHTED SEMANTIC $P@N^* = P@N^*_{\text{EXACT}} + P@N^*_{\text{SEM}}$ SEPARATES OUT THE CONTRIBUTIONS FROM EXACT AND EXCLUSIVELY SEMANTIC MATCHES. SCORES ARE GIVEN AS PERCENTAGES (%).

Model	$P@N^*$	$P@N^*_{\text{exact}}$	$P@N^*_{\text{sem}}$
VISIONSPEECHCNN	47.5	22.3	25.3
VISIONCNN	44.7	19.2	25.5
SUPERVISEDBOWCNN	50.0	38.4	11.7

considers only non-exact semantic matches). Table IV shows that the $P@N^*$ of SUPERVISEDBOWCNN is dominated by correct exact predictions, with a contribution of 11.7% from semantic matches. In contrast, VISIONSPEECHCNN makes more correct semantic (25.3%) than exact (22.3%) predictions. The visual model VISIONCNN has the highest proportion of semantic matches by a small margin. The non-exact match results track well the Spearman's ρ results in Table III.

For a qualitative comparison of VISIONSPEECHCNN and SUPERVISEDBoWCNN, we passed a set of isolated segmented spoken words through both models. Figure 3 shows t-SNE embeddings [100] of representations from a 256-dimensional bottleneck layer (used for computational reasons) added between the 3000-dimensional ReLU layer and the final output of both models. Although words are well-separated in the case of the transcription-supervised model (a), the visually grounded model are more successful in mapping semantically related spoken words to similar embeddings (b): related spoken words like ‘bike’, ‘rides’ and ‘riding’ have similar embeddings, as do ‘air’ and ‘jumps’, and ‘football’ and ‘ball’ (which is also closer to ‘soccer’).

What is the best we could do with a cascaded approach? The cheating models TEXTWUP and TEXTPARAGRAM (rows 6 and 7, Table III) represent the setting where we have access to a perfect ASR system. On the metrics that use hard ground truth labels, these two models perform best. However, VISIONCNN and VISIONSPEECHCNN perform better on Spearman’s ρ . As noted, the visually trained models are particularly strong in matching non-exact semantic keywords.

How does cascaded retrieval perform in low-resource settings in comparison to visual grounding? The previous question above considered the setting where perfect ASR is available. Our main focus is on extreme low-resource settings where no ASR is available. But in some low-resource settings, an intermediate situation may exist where limited amounts of data might still be available for ASR development. It is therefore interesting to consider how a cascaded model such as TEXTPARAGRAM would perform with actual ASR and how performance changes as ASR accuracy deteriorates. To simulate systematically worse low-resource settings, we artificially introduce errors in the ground truth transcriptions,⁵ and then apply TEXTPARAGRAM to this simulated ASR output. We also use an actual ASR system (Google Cloud Speech-to-Text⁶), but this is only a reference since in typical low-resource settings such an accurate system would not be available.

Semantic speech retrieval results are shown in Table V, with SIMASR indicating the cascaded approach with the simulated ASR systems. Performance across all metrics systematically deteriorates as ASR accuracy becomes worse. The cascaded approach with SIMASR at a word error rate (WER) of 50% performs similarly to VISIONSPEECHCNN on $P@10$ and $P@N$ (bottom row), but worse on all other metrics. This indicates that VISIONSPEECHCNN could prove useful even when ASR is available (with a relatively low accuracy). However, further experiments on truly low-resource languages are required to illustrate this more conclusively.

Are the text-based models an alternative to human judgements? Since the image-speech corpus has transcriptions, we may ask whether it is necessary to collect human annotations at all, or whether they can be generated automatically by a semantic model like our TEXT cheating models. Although the TEXT methods match hard human labels better than other

TABLE V
SEMANTIC SPEECH RETRIEVAL PERFORMANCE (%) USING A CASCADED APPROACH COMBINING ASR WITH TEXTPARAGRAM. THE WORD ERROR RATE OF EACH ASR SYSTEM IS GIVEN IN PARENTHESES. PERFECTASR REPRODUCES THE TEXTPARAGRAM RESULTS (LAST ROW) OF TABLE III. GOOGLEASR USES AN ACTUAL HIGH-RESOURCE ASR SYSTEM CASCADED WITH TEXTPARAGRAM. SIMASR USES A SIMULATED ASR SYSTEM WITH VARYING ERROR RATES CASCADED WITH TEXTPARAGRAM.

Method (WER)	$P@10$	$P@N$	EER	AP	Spear. ρ
PERFECTASR (0%)	88.8	64.0	14.3	60.1	31.6
GOOGLEASR (8.6%)	88.2	62.8	15.5	58.8	30.8
SIMASR (8.6%)	85.5	61.0	16.7	56.0	27.7
SIMASR (20%)	77.5	55.3	20.7	49.1	26.8
SIMASR (50%)	60.5	40.9	30.9	31.4	19.2
SIMASR (80%)	34.6	22.3	40.8	15.0	9.6
VISIONSPEECHCNN	58.8	39.7	23.9	39.4	32.4

approaches in Table III, they are far from perfect. Furthermore, when evaluated against the soft annotator counts, VISIONCNN and VISIONSPEECHCNN actually do better, indicating that the TEXT models cannot replace human judgements.

VII. CONCLUSIONS AND FUTURE WORK

We investigated how a model that learns from parallel images and unlabelled speech captures aspects of semantics in speech. We collected a new data set for a semantic speech retrieval task, where the aim is to retrieve utterances that are semantically relevant given a written query keyword. Without seeing any parallel speech and text, the visually grounded model achieves a semantic $P@10$ of almost 60%. Although a model trained on transcriptions is superior on some metrics, the vision-speech model retrieves more than double the number of non-verbatim semantic matches and is a better predictor of the actual soft human ratings.

Since visual context seems to provide different information from transcriptions for semantic speech retrieval, future work could consider how these two supervision signals could be combined when both are available. We would also like to explore whether our model could be used to localise the parts of the input signal giving rise to a particular (semantic) label, similarly to analyses in studies that inspired this work [7], [67]. Finally, many cognitively motivated studies have proposed to combine spoken and visual input to model infant language acquisition, but few use real unlabelled speech with natural visual input as we do here. The cognitive implications of our model could be a topic for future exploration.

We encourage interested readers to make use of our data and associated task, which we make available at: https://github.com/kamperh/semantic_flickraudio/.

ACKNOWLEDGEMENTS

We would like to thank Shane Settle, John Wieting, and Michael Roth for helpful discussions and input.

⁵Given a target word error rate e , for each word in the ground truth transcriptions, we introduce an error to it with probability e and then uniformly pick a substitution, deletion or insertion.

⁶<https://cloud.google.com/speech-to-text/>

REFERENCES

- [1] G. Adda, S. Stüker, M. Adda-Decker, O. Ambourou, L. Besacier, D. Blachon, H. Bonneau-Maynard, P. Godard, F. Hamlaoui, D. Idiatov *et al.*, “Breaking the unwritten language barrier: The BULB project,” *Proc. SLTU*, 2016.
- [2] G. Synnaeve, T. Schatz, and E. Dupoux, “Phonetics embedding learning with side information,” in *Proc. SLT*, 2014.
- [3] S. Settle, K. Levin, H. Kamper, and K. Livescu, “Query-by-example search with discriminative neural acoustic word embeddings,” in *Proc. Interspeech*, 2017.
- [4] L. Duong, A. Anastasopoulos, D. Chiang, S. Bird, and T. Cohn, “An attentional model for speech translation without transcription,” in *Proc. NAACL*, 2016, pp. 949–959.
- [5] S. Bansal, H. Kamper, A. Lopez, and S. J. Goldwater, “Towards speech-to-text translation without speech recognition,” in *Proc. EACL*, 2017.
- [6] R. J. Weiss, J. Chorowski, N. Jaitly, Y. Wu, and Z. Chen, “Sequence-to-sequence models can directly translate foreign speech,” in *Proc. Interspeech*, 2017.
- [7] D. Palaz, G. Synnaeve, and R. Collobert, “Jointly learning to locate and classify words using convolutional networks,” in *Proc. Interspeech*, 2016.
- [8] G. Chrupała, L. Gelderloos, and A. Alishahi, “Representations of language in a model of visually grounded speech signal,” *arXiv preprint arXiv:1702.01991*, 2017.
- [9] J. Luo, B. Caputo, A. Zweig, J.-H. Bach, and J. Anemüller, “Object category detection using audio-visual cues,” in *Proc. ICVS*, 2008.
- [10] V. Krunić, G. Salvi, A. Bernardino, L. Montesano, and J. Santos-Victor, “Affordance based word-to-meaning association,” in *Proc. ICRA*, 2009.
- [11] T. Taniguchi, T. Nagai, T. Nakamura, N. Iwahashi, T. Ogata, and H. Asoh, “Symbol emergence in robotics: A survey,” *Adv. Robotics*, vol. 30, no. 11-12, pp. 706–728, 2016.
- [12] G. Synnaeve, M. Versteegh, and E. Dupoux, “Learning words from images and speech,” in *NIPS Workshop Learn. Semantics*, 2014.
- [13] D. Harwath, A. Torralba, and J. R. Glass, “Unsupervised learning of spoken language with visual context,” in *Proc. NIPS*, 2016.
- [14] G. Chrupała, L. Gelderloos, and A. Alishahi, “Representations of language in a model of visually grounded speech signal,” in *Proc. ACL*, 2017.
- [15] J. Drexler and J. R. Glass, “Analysis of audio-visual features for unsupervised speech recognition,” in *Proc. GLU*, 2017.
- [16] K. Leidal, D. Harwath, and J. Glass, “Learning modality-invariant representations for speech and images,” *Proc. ASRU*, 2017.
- [17] O. Scharenborg *et al.*, “Linguistic unit discovery from multi-modal inputs in unwritten languages: Summary of the ‘Speaking Rosetta’ JSALT 2017 workshop,” *arXiv preprint arXiv:1802.05092*, 2018.
- [18] D. Harwath, A. Recasens, D. Surís, G. Chuang, A. Torralba, and J. Glass, “Jointly discovering visual objects and spoken words from raw sensory input,” *arXiv preprint arXiv:1804.01452*, 2018.
- [19] D. Harwath, G. Chuang, and J. R. Glass, “Vision as an interlingua: Learning multilingual semantic embeddings of untranscribed speech,” in *Proc. ICASSP*, 2018.
- [20] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” *Proc. Interspeech*, 2017.
- [21] C. Yu and L. B. Smith, “Rapid word learning under uncertainty via cross-situational statistics,” *Psychol. Sci.*, vol. 18, no. 5, pp. 414–420, 2007.
- [22] T. Cunillera, E. Camara, M. Laine, and A. Rodriguez-Fornells, “Speech segmentation is facilitated by visual cues,” *Q. J. Exp. Psychol.*, vol. 63, no. 2, pp. 260–274, 2010.
- [23] E. D. Thiessen, “Effects of visual information on adults and infants auditory statistical learning,” *Cognitive Sci.*, vol. 34, no. 6, pp. 1093–1106, 2010.
- [24] D. K. Roy and A. P. Pentland, “Learning words from sights and sounds: A computational model,” *Cognitive Sci.*, vol. 26, no. 1, pp. 113–146, 2002.
- [25] C. Yu and D. H. Ballard, “A multimodal learning interface for grounding spoken language in sensory perceptions,” *ACM T. Appl. Perception*, vol. 1, no. 1, pp. 57–80, 2004.
- [26] L. Ten Bosch and B. Cranen, “A computational model for unsupervised word discovery,” in *Proc. Interspeech*, 2007.
- [27] G. Aimetti, “Modelling early language acquisition skills: Towards a general statistical learning mechanism,” in *Proc. EACL*, 2009.
- [28] J. Driesen and H. Van hamme, “Modelling vocabulary acquisition, adaptation and generalization in infants using adaptive Bayesian PLSA,” *Neurocomputing*, vol. 74, no. 11, pp. 1874–1882, 2011.
- [29] O. Räsänen and H. Rasilo, “A joint model of word segmentation and meaning acquisition through cross-situational learning,” *Psychol. Rev.*, vol. 122, no. 4, pp. 792–829, 2015.
- [30] 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.
- [31] J. M. Siskind, “A computational study of cross-situational techniques for learning word-to-meaning mappings,” *Cognition*, vol. 61, no. 1, pp. 39–91, 1996.
- [32] M. C. Frank, N. D. Goodman, and J. B. Tenenbaum, “Using speakers’ referential intentions to model early cross-situational word learning,” *Psychol. Sci.*, vol. 20, no. 5, pp. 578–585, 2009.
- [33] E. V. Clark, “How language acquisition builds on cognitive development,” *Trends Cogn. Sci.*, vol. 8, no. 10, pp. 472–478, 2004.
- [34] A. Jansen, E. Dupoux, S. J. Goldwater, M. Johnson, S. Khudanpur, K. Church, N. Feldman, H. Hermansky, F. Metzke, R. Rose, M. Seltzer, P. Clark, I. McGraw, B. Varadarajan, E. Bennett, B. Borschinger, J. Chiu, E. Dunbar, A. Fourtassi, D. Harwath, C.-y. Lee, K. Levin, A. Norouziyan, V. Peddinti, R. Richardson, T. Schatz, and S. Thomas, “A summary of the 2012 JHU CLSP workshop on zero resource speech technologies and models of early language acquisition,” in *Proc. ICASSP*, 2013.
- [35] M. Versteegh, X. Anguera, A. Jansen, and E. Dupoux, “The Zero Resource Speech Challenge 2015: Proposed approaches and results,” in *Proc. SLTU*, 2016.
- [36] B. Varadarajan, S. Khudanpur, and E. Dupoux, “Unsupervised learning of acoustic sub-word units,” in *Proc. ACL*, 2008.
- [37] C.-y. Lee and J. R. Glass, “A nonparametric Bayesian approach to acoustic model discovery,” in *Proc. ACL*, 2012.
- [38] Y. Zhang, R. Salakhutdinov, H.-A. Chang, and J. R. Glass, “Resource configurable spoken query detection using deep Boltzmann machines,” in *Proc. ICASSP*, 2012.
- [39] A. S. Park and J. R. Glass, “Unsupervised pattern discovery in speech,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 16, no. 1, pp. 186–197, 2008.
- [40] C.-y. Lee, T. O’Donnell, and J. R. Glass, “Unsupervised lexicon discovery from acoustic input,” *Trans. ACL*, vol. 3, pp. 389–403, 2015.
- [41] H. Kamper, A. Jansen, and S. J. Goldwater, “Unsupervised word segmentation and lexicon discovery using acoustic word embeddings,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 24, no. 4, pp. 669–679, 2016.
- [42] H. Kamper, “Unsupervised neural and bayesian models for zero-resource speech processing,” Ph.D. dissertation, University of Edinburgh, UK, 2016.
- [43] R. Socher and L. Fei-Fei, “Connecting modalities: Semi-supervised segmentation and annotation of images using unaligned text corpora,” in *Proc. CVPR*, 2010.
- [44] J. Weston, S. Bengio, and N. Usunier, “WSABIE: Scaling up to large vocabulary image annotation,” in *Proc. IJCAI*, 2011.
- [45] M. Hodosh, P. Young, and J. Hockenmaier, “Framing image description as a ranking task: Data, models and evaluation metrics,” *J. Artif. Intell. Res.*, vol. 47, pp. 853–899, 2013.
- [46] A. Karpathy, A. Joulin, and F. F. F. Li, “Deep fragment embeddings for bidirectional image sentence mapping,” in *Proc. NIPS*, 2014.
- [47] A. Farhadi, M. Hejrati, M. A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth, “Every picture tells a story: Generating sentences from images,” in *Proc. ECCV*, 2010.
- [48] Y. Yang, C. L. Teo, H. Daumé III, and Y. Aloimonos, “Corpus-guided sentence generation of natural images,” in *Proc. EMNLP*, 2011.
- [49] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg, “Babytalk: Understanding and generating simple image descriptions,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 12, pp. 2891–2903, 2013.
- [50] R. Bernardi, R. Cakici, D. Elliott, A. Erdem, E. Erdem, N. Ikiçler-Cinbis, F. Keller, A. Muscat, and B. Plank, “Automatic description generation from images: A survey of models, datasets, and evaluation measures,” *J. Artif. Intell. Res.*, vol. 55, pp. 409–442, 2016.
- [51] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, F. Saenko, and T. Darrell, “Long-term recurrent convolutional networks for visual recognition and description,” in *Proc. CVPR*, 2015.
- [52] H. Fang, S. Gupta, F. Iandola, R. K. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. C. Platt *et al.*, “From captions to visual concepts and back,” in *Proc. CVPR*, 2015.
- [53] A. Karpathy and L. Fei-Fei, “Deep visual-semantic alignments for generating image descriptions,” in *Proc. CVPR*, 2015.
- [54] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, “Show and tell: A neural image caption generator,” in *Proc. CVPR*, 2015.

- [55] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *Proc. ICML*, 2015.
- [56] X. Chen and C. Lawrence Zitnick, "Mind's eye: A recurrent visual representation for image caption generation," in *Proc. CVPR*, 2015, pp. 2422–2431.
- [57] Y. Feng and M. Lapata, "Visual information in semantic representation," in *Proc. NAACL*. Association for Computational Linguistics, 2010, pp. 91–99.
- [58] C. Silberer and M. Lapata, "Grounded models of semantic representation," in *Proc. EMNLP*, 2012.
- [59] E. Bruni, N.-K. Tran, and M. Baroni, "Multimodal distributional semantics," *J. Artif. Intell. Res.*, vol. 49, no. 2014, pp. 1–47, 2014.
- [60] E. Bruni, G. Boleda, M. Baroni, and N.-K. Tran, "Distributional semantics in technicolor," in *Proc. ACL*, 2012.
- [61] S. Roller and S. S. Im Walde, "A multimodal LDA model integrating textual, cognitive and visual modalities," in *Proc. EMNLP*, 2013.
- [62] C. Silberer and M. Lapata, "Learning grounded meaning representations with autoencoders," in *Proc. ACL*, 2014.
- [63] F. Sun, D. Harwath, and J. R. Glass, "Look, listen, and decode: Multimodal speech recognition with images," in *Proc. SLT*, 2016.
- [64] A. Gupta, Y. Miao, L. Neves, and F. Metze, "Visual features for context-aware speech recognition," in *Proc. ICASSP*, 2017.
- [65] T. J. Hazen, B. Sherry, and M. Adler, "Speech-based annotation and retrieval of digital photographs," in *Proc. Interspeech*, 2007.
- [66] X. Anguera, J. Xu, and N. Oliver, "Multimodal photo annotation and retrieval on a mobile phone," in *Proc. ICMI*, 2008.
- [67] D. Harwath and J. R. Glass, "Learning word-like units from joint audio-visual analysis," in *Proc. ACL*, 2017.
- [68] J. Xu and W. B. Croft, "Query expansion using local and global document analysis," in *Proc. SIGIR*, 1996.
- [69] J. Graupmann, R. Schenkel, and G. Weikum, "The SphereSearch engine for unified ranked retrieval of heterogeneous XML and web documents," in *Proc. VLDB*, 2005.
- [70] G. A. Miller, "WordNet: a lexical database for english," *Commun. ACM*, vol. 38, no. 11, pp. 39–41, 1995.
- [71] F. Diaz, B. Mitra, and N. Craswell, "Query expansion with locally-trained word embeddings," in *Proc. ACL*, 2016.
- [72] D. Roy, D. Paul, M. Mitra, and U. Garain, "Using word embeddings for automatic query expansion," in *SIGIR Workshop on Neural Information Retrieval*, 2016.
- [73] J. G. Wilpon, L. R. Rabiner, C.-H. Lee, and E. Goldman, "Automatic recognition of keywords in unconstrained speech using hidden Markov models," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 38, no. 11, pp. 1870–1878, 1990.
- [74] I. Szöke, P. Schwarz, P. Matejka, L. Burget, M. Karafiát, M. Fapso, and J. Cernocký, "Comparison of keyword spotting approaches for informal continuous speech," in *Proc. Interspeech*, 2005.
- [75] A. Garcia and H. Gish, "Keyword spotting of arbitrary words using minimal speech resources," in *Proc. ICASSP*, 2006.
- [76] J. S. Garofolo, C. G. Auzanne, and E. M. Voorhees, "The TREC spoken document retrieval track: A success story," in *Content-Based Multimedia Information Access-Volume 1*, 2000, pp. 1–20.
- [77] C. Chelba, T. J. Hazen, and M. Saraclar, "Retrieval and browsing of spoken content," *IEEE Signal Proc. Mag.*, vol. 25, no. 3, 2008.
- [78] H.-Y. Lee, T.-H. Wen, and L.-S. Lee, "Improved semantic retrieval of spoken content by language models enhanced with acoustic similarity graph," in *Proc. SLT*, 2012.
- [79] 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.
- [80] Y.-C. Li, H.-y. Lee, C.-T. Chung, C.-a. Chan, and L.-s. Lee, "Towards unsupervised semantic retrieval of spoken content with query expansion based on automatically discovered acoustic patterns," in *Proc. ASRU*, 2013.
- [81] D. F. Harwath, T. J. Hazen, and J. R. Glass, "Zero resource spoken audio corpus analysis," in *Proc. ICASSP*, 2013.
- [82] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 28, no. 4, pp. 357–366, 1980.
- [83] Y. Aytar, C. Vondrick, and A. Torralba, "SoundNet: Learning sound representations from unlabeled video," in *Proc. NIPS*, 2016.
- [84] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. CVPR*, 2009.
- [85] K. Barnard, P. Duygulu, D. Forsyth, N. d. Freitas, D. M. Blei, and M. I. Jordan, "Matching words and pictures," *J. Mach. Learn. Res.*, vol. 3, pp. 1107–1135, 2003.
- [86] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid, "TagProp: Discriminative metric learning in nearest neighbor models for image auto-annotation," in *Proc. ICCV*, 2009.
- [87] M. Chen, A. X. Zheng, and K. Q. Weinberger, "Fast image tagging," in *Proc. ICML*, 2013.
- [88] 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.
- [89] 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.
- [90] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [91] D. Harwath and J. R. Glass, "Deep multimodal semantic embeddings for speech and images," in *Proc. ASRU*, 2015.
- [92] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [93] E. Agirre, E. Alfonseca, K. Hall, J. Kravalova, M. Paşca, and A. Soroa, "A study on similarity and relatedness using distributional and WordNet-based approaches," in *Proc. HLT-NAACL*, 2009.
- [94] F. Hill, R. Reichart, and A. Korhonen, "SimLex-999: Evaluating semantic models with (genuine) similarity estimation," *Comput. Linguist.*, vol. 41, no. 4, 2015.
- [95] T. J. Hazen, W. Shen, and C. White, "Query-by-example spoken term detection using phonetic posteriorgram templates," in *Proc. ASRU*, 2009.
- [96] Y. Zhang and J. R. Glass, "Unsupervised spoken keyword spotting via segmental DTW on Gaussian posteriorgrams," in *Proc. ASRU*, 2009.
- [97] Z. Wu and M. Palmer, "Verbs semantics and lexical selection," in *Proc. ACL*, 1994.
- [98] J. Wieting, M. Bansal, K. Gimpel, and K. Livescu, "From paraphrase database to compositional paraphrase model and back," *Trans. ACL*, vol. 3, pp. 345–358, 2015.
- [99] —, "Towards universal paraphrastic sentence embeddings," *Proc. ICLR*, 2016.
- [100] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. Nov, pp. 2579–2605, 2008.