Visually grounded few-shot word learning in low-resource settings

Leanne Nortje, Dan Oneață, Herman Kamper

arXiv:2306.11371v3 [eess.AS] 18 Apr 2024

***Abstract*—We propose a visually grounded speech model that learns new words and their visual depictions from just a few word-image example pairs. Given a set of test images and a spoken query, we ask the model which image depicts the query word. Previous work has simplified this few-shot learning problem by either using an artificial setting with digit word-image pairs or by using a large number of examples per class. Moreover, all previous studies were performed using English speech-image data. We propose an approach that can work on natural word-image pairs but with less examples, i.e. fewer shots, and then illustrate how this approach can be applied for multimodal few-shot learning in a real low-resource language, Yorùbá. Our approach involves using the given word-image example pairs to mine new unsupervised word-image training pairs from large collections of unlabelled speech and images. Additionally, we use a word-to-image attention mechanism to determine word-image similarity. With this new model, we achieve better performance with fewer shots than previous approaches on an existing English benchmark. Many of the model's mistakes are due to confusion between visual concepts co-occurring in similar contexts. The experiments on Yorùbá show the benefit of transferring knowledge from a multimodal model trained on a larger set of English speech-image data.**[1]

***Index Terms*—few-shot learning, multimodal modelling, visually grounded speech models, word acquisition, low-resource language**

I. Introduction

Speech recognition for low-resource languages faces a major obstacle: it requires large amounts of transcribed data for development [1]. In some extreme cases, it might even be impossible to get any labelled data, e.g. when dealing with an unwritten language. This is in stark contrast to infants that learn words without access to any transcriptions [2]–[6]. This is one motivation for recent studies into multimodal few-shot learning [7]–[9]: learning new concepts from a few examples, where each example consists of instances of the same concept but from different modalities. E.g., imagine a robot seeing a picture of a *zebra*, *kite* and *sheep* while also hearing the spoken word for each concept. After seeing this small set of examples (called a support set) the robot is prompted to identify which image in an unseen set corresponds to the word "zebra".

Leanne Nortje and Herman Kamper are with the Department of Electrical and Electronic Engineering, Stellenbosch University, Stellenbosch, 7600, South Africa (email: nortjeleanne@gmail.com; kamperh@sun.ac.za).

Dan Oneață is with the University Politehnica of Bucharest, Bucharest, RO-060042, Romania (e-mail: dan.oneata@gmail.com).

Google DeepMind funded Leanne Nortje. Dan Oneață was supported by European Union's HORIZON-CL4-2021-HUMAN-01 research and innovation program under grant agreement No. 101070190 AI4Trust. We would like to thank Tyler Miller and David Harwath for helping with the few-shot retrieval comparisons, and Benjamin van Niekerk for helping with QbERT.

[1]Code and models: Project webpage.

Building off of a growing number of studies in visually grounded speech modelling [10]–[16], we consider this multimodal problem of learning the spoken form of a word and its visual depiction from only a few paired word-image examples. Multimodal few-shot speech-image learning was first introduced in [7] and then extended in [8] and [9]. But these studies were performed in an artificial setting where spoken isolated digits were paired with MNIST images of digits. This shortcoming was recently addressed by Miller and Harwath [17], who considered multimodal few-shot learning on isolated words paired with natural images. Their specific focus was on learning a new concept while not forgetting previously learned concepts, i.e. dealing with the problem of catastrophic forgetting. (We do not explicitly focus on the catastrophic forgetting problem here, although we do evaluate using the same setup as [17].) While their methods performed well in a few-shot retrieval task with five classes, they required a relatively large number of samples per class, i.e. many "shots". Our first overarching aim is to do visually grounded multimodal few-shot learning on natural images with fewer shots. All previous studies also performed experiments using English speech-image data. Our second goal, therefore, is to present a few-shot evaluation on a real low-resource language.

Our new multimodal few-shot approach combines two core ideas. Firstly, we use the support set to "mine" new noisy word-image pairs from unlabelled speech and image collections. Concretely, each spoken word example in the support set is compared to each utterance in an unlabelled speech corpus; we use a new query-by-example approach to identify segments in the search utterances that match the word in the support set. We follow a similar approach for mining additional images from the few-shot classes by using cosine distance between pretrained image embeddings. The mined words and images are then paired up, thereby artificially increasing the size of our support set in an unsupervised way. This mining scheme is very similar to that followed in [9], where it was used on digit image-speech data with simpler within-modality comparisons. Secondly, our new approach is based on a model with a word-to-image attention mechanism. This multimodal attention network (MattNet), takes a single word embedding and calculates its correspondence to each pixel embedding to learn how the word is depicted within an image. This is similar to the vision attention part of the model from [18], where the goal was to localise visual keywords in speech (not in a few-shot setting).

Using the English SpokenCOCO speech-image dataset [19], two evaluation settings are considered. We first evaluate our approach on the few-shot retrieval task also used in [17]. We show that MattNet achieves higher retrieval scores for fewer

shots than [17]'s models. Secondly, we evaluate our approach in a more conventional few-shot classification task where it only needs to correctly distinguish between classes seen in the support set. We show that we can achieve five-way accuracies higher than 80% with as little as five shots.

Still using the English SpokenCOCO models, we then perform an exhaustive analysis to understand the parts of our approach that are most essential, the characteristics of mistakes, and how performance differs across keywords. This includes qualitative results showing that the few-shot model can localise objects in images given a spoken word query.

To address our second overarching goal, we finally do multimodal few-shot learning using a Yorùbá speech-image dataset [20], illustrating for the first time that these few-shot approaches can be applied to a real low-resource language.

This work is an extension of the conference paper [21], where the main results on the SpokenCOCO benchmark were presented (a slightly improved model is used here). The current paper extends this work with a thorough analysis, ablation experiments, and the application of the model to a low-resource language. To summarise, we make the following contributions: (1) we introduce a new mining scheme operating on natural images and speech, (2) we introduce a new attention-based model for multimodal few-shot learning, (3) we give a thorough analysis of the proposed approach, and (4) apply the approach in a new low-resource multimodal few-shot learning benchmark.

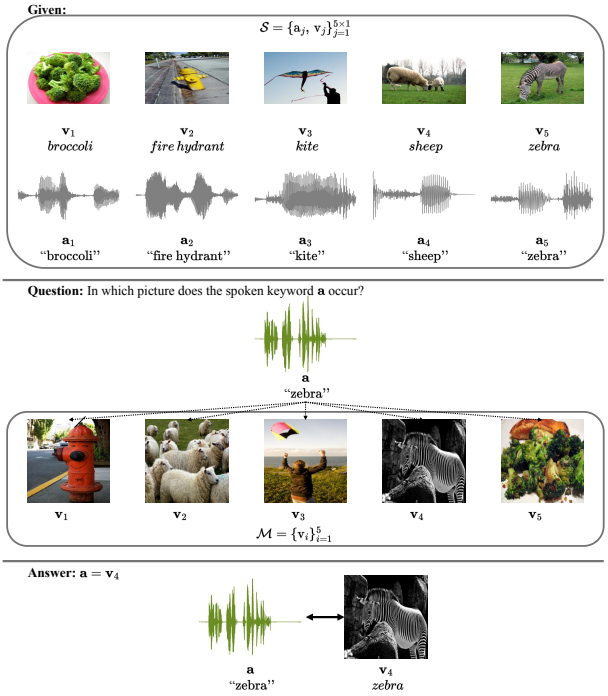

Fig. 1. Given the few examples in the support set $\mathcal{S}$, the multimodal few-shot classification task is to e.g. identify the image depicting the word "zebra" from a set of unseen images.

II. Visually grounded few-shot learning and evaluation

We train a model on a few spoken word-image examples. The set of K examples per class is called the support set $\mathcal{S}$. Each pair in $\mathcal{S}$ consists of an isolated spoken word $\mathbf{a}_j$ and a corresponding image $\mathbf{v}_j$. For the *one-shot* case shown in the top part of Fig. 1, $\mathcal{S}$ consists of one word-image example pair for each of the L classes. For the L-way K-shot task, the support set $\mathcal{S} = \{\mathbf{a}_j, \mathbf{v}_j\}_{j=1}^{L\times K}$ contains K word-image example pairs for each of the L classes. In this work, we use a few-shot model for two tasks, as we describe next.

A. Visually grounded few-shot word classification

In this task, illustrated in the middle and bottom of Fig. 1, we are given an unseen isolated spoken word query $\mathbf{a}$ and prompted to identify the corresponding image in a matching set $\mathcal{M} = \{\mathbf{v}_i\}_{i=1}^{L}$ of unseen test images. $\mathcal{M}$ contains one image depicting each of the L classes. Neither the test-time speech query $\mathbf{a}$ nor any images in $\mathcal{M}$ occur in the support set. This image-speech task was considered in [7]–[9], but here, for the first time, we use natural images instead of isolated digit images. In contrast to the task described next, this is conventional few-shot classification where the model only needs to correctly distinguish between classes seen in the support set, i.e. there are no other background or imposter classes.

B. Visually grounded few-shot retrieval

In contrast, in this task the goal is to test whether a model can search through a large collection of images and retrieve those that depict a few-shot speech query, i.e. the matching set $\mathcal{M}$ in this case contains images that depict the L few-shot classes but also images that depict other classes. These additional images might contain completely unseen classes, or background classes potentially seen during pretraining of the few-shot model. The model is penalised if it retrieves one of these imposter images. This few-shot retrieval task was proposed in [17]. Their interest was specifically in measuring catastrophic forgetting. Since their task requires a model to distinguish between few-shot classes and other classes, it can be used to not only determine whether models can be updated to learn new classes from only a few examples, but also how well the model remembers previously learned (background) classes. We do not explicitly focus on the catastrophic forgetting problem, but we want to compare to [17]. Therefore, we also consider this retrieval task.

For both tasks we need a distance metric $D_{\mathcal{S}}(\mathbf{a}, \mathbf{v})$ between instances from the speech and vision modalities. Next we describe the model that we use to compute this distance metric.

III. Multimodal few-shot attention

Our approach for determining $D_{\mathcal{S}}(\mathbf{a}, \mathbf{v})$ relies on two core components: a model with a word-to-image attention mechanism and a method to mine pairs using a few ground truth word-image examples (given in the support set).

A. Word-to-image attention mechanism

Our model is shown in Fig. 2 and we call it MattNet (Multimodal ATTention NETwork). We adapt the multimodal localising attention model of [18] that consists of an audio and a vision branch. For the vision branch, we replace ResNet50 [22]

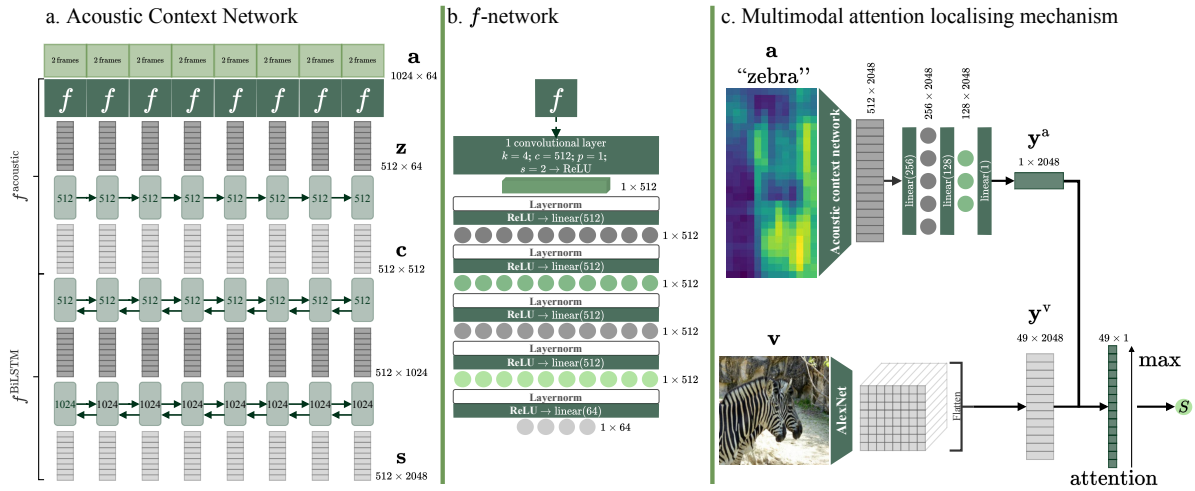

Fig. 2. MATTNET consists of (c) a vision and an audio network. The audio network consists of (a + b) an acoustic context network and a BiLSTM network. The audio and vision networks are connected with a word-to-image attention mechanism.

with an adaption of AlexNet [23] to encode an image input $\mathbf{v}$ into a sequence of embeddings $\mathbf{y}^{\text{v}}$. Originally [21], additional linear layers were used after the image embeddings, but removing these did not impact performance. For the audio branch, we use the same audio subnetwork as [18] that consists of an acoustic network f^{acoustic} which extracts speech features from a spoken input $\mathbf{a}$. However, [18] takes an entire spoken utterance as $\mathbf{a}$, whereas we use a single isolated spoken word. We also add a few linear layers to the BiLSTM network f^{BiLSTM} to encode the speech features into a single audio embedding $\mathbf{y}^{\text{a}}$, similar to acoustic word embeddings [24]–[27]. We connect the vision and audio branches with a multimodal attention mechanism to compare the word embedding $\mathbf{y}^{\text{a}}$ to each embedding in $\mathbf{y}^{\text{v}}$.

To get this word-to-image attention mechanism, we take the keyword localising attention mechanism of [18] which detects whether certain keywords occur in both spoken utterances and images. However, we aim to only detect whether a single isolated spoken word occurs somewhere within an image. More specifically, we calculate attention weights over the image embeddings by calculating the dot product between $\mathbf{y}^{\text{a}}$ and each embedding in $\mathbf{y}^{\text{v}}$. By taking the maximum over the attention scores, we get a similarity score S. The higher S, the more probable it is that the spoken word corresponds to one or more objects in the image. If S is low, it is less probable that any object in the image corresponds to the spoken word.

We train MATTNET with a contrastive loss:

$$l = \text{MSE}\left(S(\mathbf{a}, \mathbf{v}), 100\right) + \sum_{i=1}^{N_{\text{pos}}} \text{MSE}\left(\left[S(\mathbf{a}, \mathbf{v}_i^+), S(\mathbf{a}_i^+, \mathbf{v})\right], 100\right) + \sum_{i=1}^{N_{\text{neg}}} \text{MSE}\left(\left[S\left(\mathbf{a}_i^-, \mathbf{v}\right), S(\mathbf{a}, \mathbf{v}_i^-), S(\mathbf{a}, \mathbf{v}_i^{\text{bg}})\right], 0\right), \tag{1}$$

where S is calculated with MATTNET and we limit $S \in [0, 100]$. Intuitively this loss should push $\mathbf{a}$, $\mathbf{v}$ and the positive examples $\mathbf{a}_i^+$ and $\mathbf{v}_i^+$ closer together using a mean square error (MSE) that pushes the list of similarities to 100. At the same time the loss should push the negative examples $\mathbf{a}_i^-$, $\mathbf{v}_i^-$ and $\mathbf{v}_i^{\text{bg}}$ away from these positives (through the MSE to 0). But before we can do this, we need positive ($\mathbf{a}_i^+$, $\mathbf{v}_i^+$), negative ($\mathbf{a}_i^-$, $\mathbf{v}_i^-$, $\mathbf{v}_i^{\text{bg}}$) and anchor ($\mathbf{a}$, $\mathbf{v}$) pairs.

B. Few-shot pair mining

For few-shot training, we only have the small number of ground truth examples in the support set $\mathcal{S}$. This would not be sufficient to train the model. To overcome this, [9] proposed a pair mining scheme: use the audio examples in $\mathcal{S}$ and compare each example to each utterance in a large collection of unlabelled audio utterances, and similarly for the images. The mined items can then be used to construct more word-image pairs for training. While in [9] the unlabelled collection of audio consisted of isolated spoken words (which was artificially segmented), here we consider an unlabelled collection of audio consisting of full spoken utterances (a more realistic scenario).

The simple isolated-word comparison approach used in [9] is not adequate for this setting. We employ another approach. We have a spoken word in our support set that we want to match to unlabelled unsegmented utterances in a large audio collection. This is similar to fuzzy string search, i.e. finding a set of strings that approximately match a given pattern. However, algorithms from string search are not directly applicable to speech since they operate on a discrete alphabet. We therefore use QbERT (query-by-example with HuBERT). The idea is to encode speech as a set of discrete units that approximate phones. Then we can apply standard string search algorithms to find examples that match a given query word. We use HuBERT [28] to map input speech into discrete units. Concretely, we use layer seven of HuBERT-Base for K-means quantisation with 100 clusters. Then we divide the units into variable-duration phone-like segments following [29]. Finally, we search the dataset by aligning the query to each utterance using the Needleman-Wunsch algorithm [30]. An alternative to QbERT would have been to use dynamic time warping (DTW), as is done in [9]. However, in a developmental experiment we found that DTW achieves an isolated word retrieval F_1 score of 76.8% while QbERT achieves 98.7%.

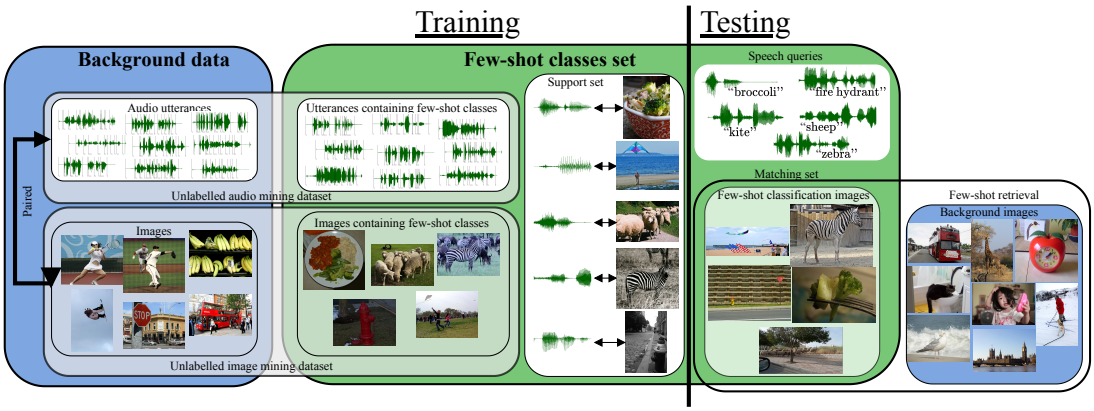

Fig. 3. The SpokenCOCO data splits used to train and evaluate the MATTNET model. The background data (blue background) consists of spoken audio utterances and images belonging to concepts not present in the support set. The mining splits consist of single-modality audio and image samples from which the training data is artificially extended, and include both background samples and samples belonging to the few-shot classes (green).

Using QbERT, we compare each spoken utterance in an unlabelled collection of audio utterances to each spoken word example in $\mathcal{S}$. For each utterance, we take the highest score across the K word examples per class and rank the utterances from highest to lowest for each class. The first n utterances with the highest scores for a class are predicted to contain the spoken form of the word. Additionally, we use QbERT's predicted word segments to isolate matched words. To mine image pairs, we follow the same steps, but instead we use AlexNet [23] to extract a single embedding for each image and use cosine distance to compare image embeddings to one another. To get word-image pairs, we mine an image from the same predicted class as a segmented word. Negative pairs are taken from the positive pairs of other classes. We also mine an extra negative image $\mathbf{v}_i^{\text{bg}}$ from a set known to not contain any of the few-shot classes (referred to as the background data, see below). Therefore, during the few-shot retrieval task, images containing few-shot classes can be distinguished from images that depicts none of the few-shot classes.

IV. ENGLISH FEW-SHOT SPEECH-IMAGE EXPERIMENTS

Before we get to an actual low-resource setting, we do experiments and perform analyses on an English benchmark.

A. Experimental setup

1) Data: For our English experiments, we use the SpokenCOCO Corpus [19] which consists of the MSCOCO [31] images with recorded spoken captions corresponding to the MSCOCO textual captions. Fig. 3 shows how we partition the dataset. Firstly, we use the setup of [17] to divide the SpokenCOCO dataset into a few-shot set (green) and a background set not containing any of the few-shot classes (blue).[2] We use the same few-shot classes as [17]: *brocolli*, *fire hydrant*, *kite*, *sheep* and *zebra*. The background data is used to pretrain MATTNET(§IV-A2).

For the few-shot set, we further divide it into training and testing according to the splits used by [17]. We use this testing set to sample isolated spoken word queries and matching images for testing. For the few-shot classification task (§II-A), we sample images only from the few-shot test set. Since the few-shot retrieval task (§II-B) requires a single large image matching set, we take all of the images in the few-shot test set as well as the images in the background test set. We sample the support set $\mathcal{S}$ from the few-shot training set (§II), using the Montreal forced aligner [32] to isolate the few-shot words. To mine pairs (§III-B), we need both an unlabelled audio and image dataset to mine pairs from; for this we use the remainder of the few-shot training data that does not include the support set as well as the background training data, shown in the left half of Fig. 3. From these unlabelled collections, we mine pairs: the $n = 600$ highest ranking examples per class (§III-B). Lastly, these pairs are split into training and validation pairs.

Utterances are parametrised as mel-spectograms with a hop length of 10 ms, a window of 25 ms and 40 mel bins. These are truncated or zero-padded to 1024 frames. Images are resized to 224×224 pixels and normalised with means and variances calculated on ImageNet [33].

2) Models: Fig. 2 illustrates our model, MATTNET (§III-A). For the image branch, we use an adaption of AlexNet [23] to get image embeddings. This image branch is also initialised using the pretrained convolutional encoder (before the classification network) of AlexNet. We use an adaption of [18]'s audio network for the audio branch. This acoustic network is pretrained on LibriSpeech [34] and the multilingual (English and Hindi) Places dataset [35] using a self-supervised contrastive predictive coding task [36]. While there are more modern alternatives for the vision and audio networks, we chose these particular variants to limit computational requirements. After initialisation, the combined MATTNET model is pretrained on the background data (blue, Fig. 3) using the contrastive speech-image retrieval loss of [37]. This is then the starting point for the model that we update using mining.

During training on the SpokenCOCO mined pairs, we take $N_{\text{pos}} = 5$ and $N_{\text{neg}} = 11$ in (1). These values were fine-tuned on the validation pairs. We train all models with Adam [38] for 100 epochs using a validation task for early stopping. For the validation task, we use the validation set to get one positive

[2] [17] refers to classes occurring in the background data as base classes.

TABLE I
THE SETUP FOR EACH MATTNET VARIANT.

Model	Pretrain	Mine training pairs	Fine-tune
MATTNET	✓	✓	✓
MATTNET, no mining	✓	✗	✓
MATTNET, no fine-tuned	✓	✗	✗

image $\mathbf{v}^+$ and one negative image $\mathbf{v}^-$ for each validation word-image $(\mathbf{a}, \mathbf{v})$ pair. The validation task measures whether the model will place $\mathbf{v}$ and $\mathbf{v}^+$ closer to $\mathbf{a}$ than it would $\mathbf{v}^-$.

To understand the models we discuss in the following section, we give the model names and their setup in Table I. Our full model is referred to as "MATTNET"; it is pretrained on background data and then fine-tuned on the mined pairs. For the "MATTNET, no mining" model, instead of training on the mined pairs, we only update the model on the samples in the support set. The "MATTNET, no fine-tuned" model consists of only the pretrained model; it is not updated with the few-shot classes in any way. To do the few-shot tasks with this model, we use the indirect few-shot method of [7], [8]: each $\mathbf{a}$ is compared to each $\mathbf{a}_j$ in $\mathcal{S}$ to find the audio example closest to the query. The image $\mathbf{v}_j$ corresponding to the closest $\mathbf{a}_j$ is then used to calculate the similarity to each image $\mathbf{v}_i$ in $\mathcal{M}$. This model can thus be seen as evaluating the quality of the embedding spaces obtained through pertaining when data from unseen classes are presented to the model.

3) Few-shot evaluation tasks: We evaluate our approach on two tasks (as explained in §II): a traditional few-shot classification task and a few-shot retrieval task. For both tasks, the K-shot L-way support set $\mathcal{S}$ contains K ground truth spoken word-image pairs for each of the $L = 5$ classes and is used to mine pairs for training and validation. In the few-shot classification task, we sample 1000 episodes where each episode contains L spoken word queries $\mathbf{a}$, one for each class, and a matching set $\mathcal{M}$ which contains one image $\mathbf{v}_i$ for each class. In the few-shot retrieval task, instead of having one image per class, $\mathcal{M}$ consists of 5000 images $\mathbf{v}_i$ where some depict a few-shot class and others do not. Here, 20 query words are taken per class and averaged to get $\mathbf{a}$. For each of the L queries $\mathbf{a}$, these 5000 images are ranked from highest to lowest similarity. The precision at N ($P@N$) score is the proportion of images in the top N highest ranking images that are from the same class as $\mathbf{a}$. N is the actual number of images in $\mathcal{M}$ that depicts the word class.

B. Experimental results

We start by comparing to two of [17]'s models on the few-shot retrieval task. The first is their naive model, which is pretrained on background classes and then fine-tuned on $K = 100$ examples for each of the $L = 5$ classes. The second is an oracle masking model in which the contrastive loss used during fine-tuning ensures that a negative image does not contain any instance of the anchor few-shot class. The results are given in Table II.[3] (Not all settings considered here were evaluated in [17], so these are indicated with dashes.)

[3]These scores are slightly different from those in [21] because we do not apply linear layers over the image encodings (§III-A).

Our full MATTNET model outperforms the oracle model across all values of K. Neither MATTNET nor the oracle masking works as well as the naive fine-tuned approach for a high number of shots (line 1, $K = 100$). In the forth line we see that our no mining approach, which is equivalent to naive fine-tuning, does worse than the naive model from [17] at $K = 100$. It is important to note that we use a different architecture. We can, however, conclude that direct fine-tuning only works with a large number of shots; as the number of shots increases, we get closer to a standard supervised learning setting, and it is therefore unsurprising that at some point naive fine-tuning starts to outperform few-shot methods. But, taking all this together, it is clear that our approach outperforms the existing methods with fewer shots.

To determine the contribution of both mining and fine-tuning, we do an experiment where we do not update MATTNET on the few-shot classes after pretraining it on the background data ("MATTNET, no fine-tuned"). We see that these two components improve the scores by roughly 20% in absolute performance when comparing lines 3 and 5 in Table II.

To further analyse the performance gains from mining, we now consider the conventional few-shot word classification task (§II-A). This task wasn't used in [17]. Table III shows that the few-shot classification scores increase as K increases when we use mined pairs. For the no fine-tuning method, the scores are lower than that of MATTNET and decrease slightly as K increases. The no mining approach in which we only update the model on the support set samples has a steep increase in scores from $K = 5$ to $K = 10$, after which the scores overtake MATTNET. Again, this makes sense as the amount of training data (K) increases. Again, these results illustrate the effectiveness of MATTNET when we have fewer shots.

Altogether, we set a competitive multimodal baseline for both few-shot retrieval and word classification in settings where the number of shots is small. However, the results do raise some questions, including the following: why do the retrieval

TABLE II
$P@N$ FEW-SHOT RETRIEVAL SCORES (%) ON THE FIVE FEW-SHOT CLASSES. K IS THE NUMBER OF SUPPORT-SET EXAMPLES PER CLASS.

Model	K			
	5	10	50	100
Naive fine-tuned [17]	–	–	–	**52.5**
Oracle masking [17]	–	8.4±0.0	24.0±0.1	35.5±0.2
MATTNET	**40.3±0.1**	**44.2±0.1**	**41.7±0.2**	43.7±0.1
MATTNET, no mining	13.2±0.6	34.8±0.7	40.9±0.3	40.5±0.5
MATTNET, no fine-tuned	22.0±0.4	24.1±0.8	22.7±0.5	23.2±1.1

TABLE III
FEW-SHOT WORD CLASSIFICATION ACCURACY (%) WHEN VARYING THE NUMBER OF SHOTS PER CLASS K.

Model	K			
	5	10	50	100
MATTNET	**80.1**	**81.1**	88.5	93.2
MATTNET, no mining	53.1	79.4	**93.3**	**95.5**
MATTNET, no fine-tuned	50.4	48.0	48.5	47.7

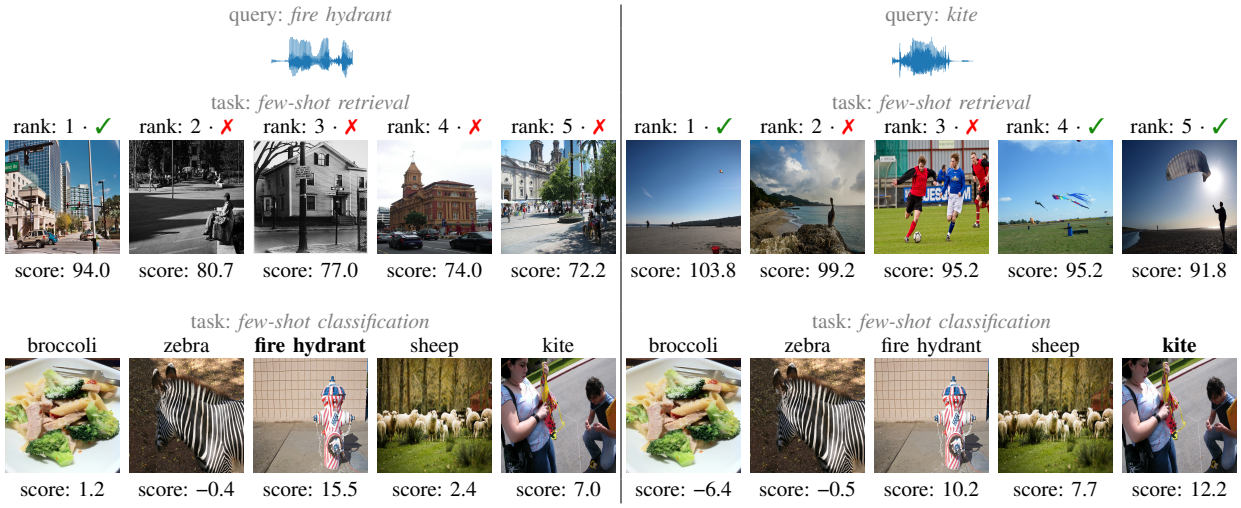

Fig. 4. Examples of retrieval and few-shot classification for two queries using the $K = 100$ MATTNET model. Concepts that associate strongly with context, such as *fire hydrant* which often appears in urban environments, are more challenging to retrieve than to classify.

scores for MATTNET plateau (Table II) but the classification scores increase (Table III) as K increases? We unpack this as part of the analysis in the next section.

V. FURTHER ANALYSIS OF ENGLISH

In this section we investigate the retrieval performance plateau we encountered in the previous section. We also present a finer-grained analysis looking into various aspects that contribute to our approach's performance.

A. Qualitative error analysis

We start with a qualitative analysis to compare the few-shot classification and retrieval tasks. Fig. 4 shows the five matching set images for classification and the five images that the model retrieves when given the spoken queries "fire hydrant" and "kite". The analysis suggests that for some classes the model depends on the contextual information to identify the class. Since the few-shot image classes and their contexts are quite distinct from one another, it does not have a significant influence on the classification task (bottom, Fig. 4). However, this has a more significant effect on retrieval since contextual information may overlap more between the few-shot images and background images (top). E.g. there might be multiple images containing streets, but only some of them contain *fire hydrants*. For classification, this learned association actually helps since none of the other few-shot classes involves streets.

B. Per-keyword analysis

Table IV presents the individual retrieval performance for each of the five few-shot keywords. (In §IV-B, specifically Table II, retrieval scores were aggregated over the keyword types.) We observe a large variance in performance across the concepts: *fire hydrant* is the most challenging keyword to retrieve (with a performance as low as 4.8% in $P@N$), while *zebra* is the easiest (over 85% when $K = 10$). The differences between keywords could be due to a number of factors, object size and word frequency probably being among the most important ones, as suggested by [17, Table 4].

Interestingly, if we look at how the performance varies with the number of shots K, we get a more nuanced picture than the one provided by the aggregated results (in Table II). While we previously observed the aggregated performance staying approximately the same with larger K, here we instead see that there are variations in both directions depending on the keyword: the retrieval scores improve with K for *sheep* and *broccoli*, but degrade for *fire hydrant* and fluctuates for *kite* and *zebra*. We speculate that as K increases for classes that have low retrieval scores (*fire hydrant*, *kite*), the model becomes more dependent on the recurring contextual (background) features, which are especially difficult to disentangle for the objects that occupy only a small part of the image. I.e. when the few-shot object is small, the model struggles to focus on the few-shot object and therefore depends on the contextual data to learn the class. Therefore, increasing K does not help. We have already seen this to some degree in the examples of Fig. 4; next we investigate this by analysing the model's localisations.

C. Localisation visualisation

Where does MATTNET focus in an image when given an audio query? We make use of the implicit localisation capabilities provided by our model. More precisely, since MATTNET projects

TABLE IV
$P@N$ FEW-SHOT RETRIEVAL SCORES (%) FOR EACH OF THE FIVE FEW-SHOT CLASSES. K IS THE NUMBER OF SUPPORT-SET EXAMPLES PER CLASS. THE NUMBER OF EXAMPLES PRESENT IN THE MATCHING SET $\mathcal{M}$ FOR EACH CLASS, N, IS GIVEN IN BRACKETS.

K	*broccoli* (57)	*fire hydrant* (62)	*kite* (91)	*sheep* (63)	*zebra* (90)
5	40.4	11.3	27.5	33.3	77.8
10	49.1	9.7	30.8	33.3	85.6
50	50.9	8.1	20.9	36.5	82.2
100	52.6	4.8	33.0	38.1	80.0

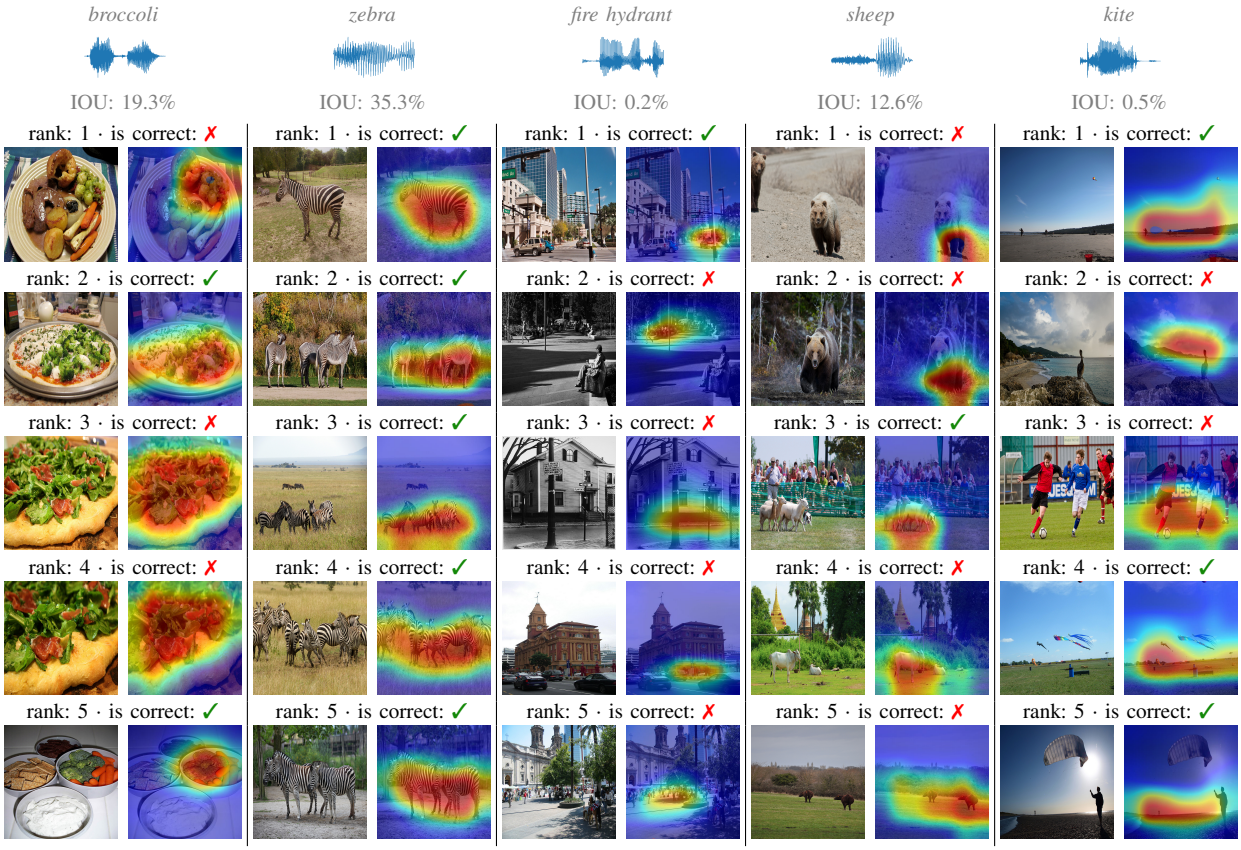

Fig. 5. The top five ranked samples for audio query corresponding to each of the five concepts using the $K = 100$ MATTNET model. For each image, we show whether it is correct (if it contains the query concept) and the attention explanation (red indicating the input regions relevant for the given audio query). The IOU values give the quantitative localisation performance of the attention explanations: the intersection over union of the binarised attentions with the ground truth annotations averaged over all images that contain a given concept.

the image features to attention scores (based on the audio feature, §III-A), we first reshape the 49-dimensional attention vector to a 7×7 matrix, which we then resize with bilinear interpolation to the original image size. The resulting map reveals the importance of the input regions: the closer the attention score is to 100, the more likely it is relevant to the input audio query. Fig. 5 shows the top five retrieved images together with their attention maps for each of the five keywords.

These visualisations are useful for performing implicit localisation of the spoken concepts in the input images, but, perhaps more importantly, they allow us to better understand the model. We observe that for "broccoli" the model selects the green food in an image; the "fire hydrant" queries are associated with urban scenes (streets and cars); and the keyword "kite" is linked to the seaside or fields. While all these associations are useful proxies for identifying the spoken keywords, they also indicate that the model tends to learn correlations with the context or other spurious features. A quantitative analysis in terms of the intersection over union (IOU) with ground truth annotations supports these findings, as shown at the top of Fig. 5. This effect is not surprising given that the model learns in a weakly-supervised manner (the training is done on full images with no explicit localisation information) and with noisy data (the mined samples are not always accurate).

With the analyses up to this point we can explain the discrepancies seen at the end of §IV. Concretely, we conclude

TABLE V
PRECISION (%) OF THE AUDIO AND IMAGE PAIRS MINED FROM THE K FEW-SHOT EXAMPLES IN $\mathcal{S}$.

K	Audio pairs	Image pairs
5	81.1	43.9
10	83.0	47.5
50	85.4	48.0
100	87.6	51.5

that the plateau seen in the retrieval scores (Table II) is due to the model associating contextual information or other spurious features with a few-shot class. E.g. associating anything green with *broccoli*, urban areas or vertical pole-like objects with *fire hydrant*, and people in fields with *kite*. This aids few-shot classification since the context or faulty features might help to distinguish the few-shot classes. However, this hurts few-shot retrieval since the model will return green objects (*field* or *grass*) when prompted with "broccoli".

D. Accuracy of the mining pairs

In §III-B we describe the method we use to mine training pairs from only the few ground truth word-image pairs in the support set $\mathcal{S}$: to find audio matches for each spoken keyword in $\mathcal{S}$ from a large collection of unlabelled speech, we use QbERT to extract possible word matches. Similarly, we use

AlexNet to find image matches for each image class in $\mathcal{S}$. This means that the audio and image pairs we use for training are not 100% correct. In order to evaluate the effect that the mined pairs have on MATTNET's performance, we report the precision of these pairs in Table V. For both the audio and image pairs, the precision increases as K increases.

E. Importance of the background data

Ideally, the usage of background negative images should force the model to throw away more contextual information and focus more on learning the few-shot objects, since the background negatives might have similar contexts to some of the few-shot cases. But since we see that contextual information is not completely ignored in the above analyses, we do the analysis in Tables VI and VII: In the second line of both tables, we have a model which is not pretrained at all and also uses no image negatives during fine-tuning on the mined pairs (§III-B), i.e. background data (§IV-A1) is not used at all. By comparing lines 1 and 2 of Table VI, we see that retrieval scores drops substantially. We can therefore conclude that adding the background information does remove some contextual information, but not all of it.

Interestingly, we see that the classification scores also decrease (lines 1 and 2 of Table VII) when we remove all background data. But this is mainly due to the effect of pretraining. To see this, we look at the contribution of the background image negatives: we train a MATTNET using the pretrained network but leave out the negative background images when fine-tuning. This model is listed in the third row of the two tables. For classification (Table VII), we see that this approach actually improves performance over MATTNET (line 1). This makes sense since the model is now fine-tuned exclusively on (mined) few-shot classes. But it also illustrates that pretraining on background data is essential. For retrieval (Table VI), we see an expected drop in performance comparing lines 1 and 3, because the latter model is not trained to distinguish between few-shot and non-few-shot classes.

TABLE VI
FEW-SHOT RETRIEVAL SCORES (%) OBTAINED FROM MATTNET TRAINED WITH AND WITHOUT NEGATIVE BACKGROUND IMAGES.

Model	K			
	5	10	50	100
MATTNET	**40.3±0.1**	**44.2±0.1**	**41.7±0.2**	**43.7±0.1**
MATTNET, no background data	18.1±1.0	23.6±1.5	26.1±0.2	23.1±0.5
MATTNET, no background image negatives	29.9±0.1	31.4±0.2	32.2±0.2	32.1±0.1

TABLE VII
FEW-SHOT CLASSIFICATION ACCURACY (%) OBTAINED FROM MATTNET TRAINED WITH AND WITHOUT NEGATIVE BACKGROUND IMAGES.

Model	K			
	5	10	50	100
MATTNET	80.1	81.1	88.5	93.2
MATTNET, no background data	65.1	64.7	75.3	77.5
MATTNET, no background image negatives	**88.0**	**90.1**	**94.8**	**95.3**

F. Adding more keywords

Can we use our model to deal with more than five classes? We leave an exhaustive investigation of this question for future work, but present some initial experiments here. Concretely, we do a test on 40 classes, which we manually select from the dataset. When training MATTNET on the 40 classes, performance on the same few-shot retrieval task for the original five classes drops marginally from 40.3% to 37.1%. This shows that, despite being trained on more classes, the model still retains most of its retrieval performance. Similarly, classification performance on the original five classes drops from 80.1% to 74.3%. When doing 40-way few-shot classification, we achieve a performance of 23.8% (this is a much more difficult task than the five-way setting).

To see what happens on the per-keyword level when MATTNET is trained and tested on 40 few-shot classes, Fig. 6 shows the individual scores. With 40 classes, there is more room for error. We also see a larger distribution of few-shot classes that the model struggles to learn. Future work will look into improving the mined image pairs specifically (since the audio pairs are more accurate).

VI. ACTUAL LOW-RESOURCE FEW-SHOT SPEECH-IMAGE EXPERIMENTS

Our ultimate goal is to do multimodal few-shot word acquisition on actual low-resource languages. To showcase our model's capabilities in this regard, we perform multimodal few-shot experiments on Yorùbá, a low-resource language spoken in Nigeria by roughly 44 million people. We apply MATTNET (§III) to the Yorùbá Flickr Audio Caption Corpus (YFACC) [20]. Because the dataset is small, we consider only the case with five shots per class ($K = 5$). We also train an English version of our model using a similarly sized dataset for comparison.

A. Experimental setup

1) Data: YFACC is an extension of the original English Flickr image-text captioning corpus [39], [40]. Concretely, for each of the 8k Flickr images, one English text caption

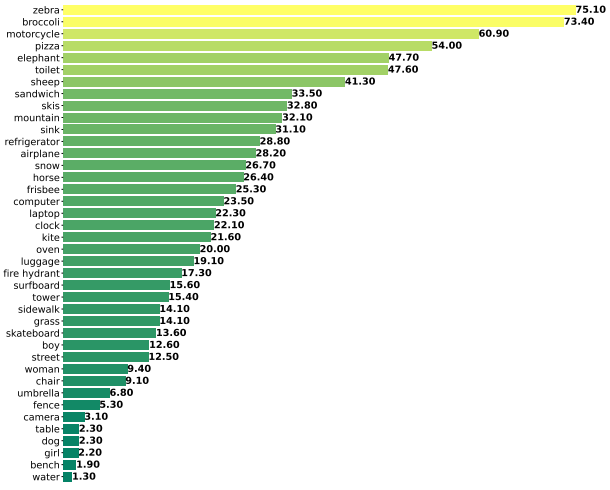

Fig. 6. The per-keyword classification scores for the 40 few-shot classes.

was translated into Yorùbá and then recorded by a single speaker. The YFACC paper [20] looked specifically at a keyword spotting task for 67 Yorùbá keywords (matching the English keywords from [41]). The YFACC test set therefore contains 500 spoken captions with manual alignments where each caption contains at least one of their 67 keywords. Since not all our original few-shot classes occur in this set (§IV-A), we choose five new classes from their 67 keywords. The Yorùbá keywords with their English translations are given in Table VIII.

Since only the test set contains alignments for the keywords, we have to sample the support set from the test set and use the remaining examples to sample episodes for the few-shot classification task. For the support set images, we crop the images to only contain a single few-shot class since the new classes frequently co-occur within images (unlike the setting for the English experiments in §IV). We stress that none of the examples in the support set occurs in the few-shot episodes' matching sets or queries. We sample the episodes for the new Flickr keywords in the same manner as in §IV-A3.

As a baseline, we also train an English system on very similar data. For this, we use the English Flickr Audio Captions Corpus (FACC) [42] which preceded and inspired YFACC. We follow the same setup as for the Yorùbá model, but replace the Yorùbá utterances from YFACC with the corresponding English utterances from FACC.

2) Models: Using the sampled support set, we mine image and audio pairs in the same way as in §V-D, with the only difference that we take the $n = 100$ highest ranking examples per class because the dataset is smaller. For mining audio pairs in Yorùbá, we still use the QbERT approach, based on an English-trained HuBERT model (§III-B), even though we are searching through unlabelled Yorùbá audio. I.e. we apply QbERT cross-lingually. For the English model, we do not use weights pretrained on the MSCOCO background data as we did before, since this data might contain paired instances of our few-shot classes. There is therefore no background speech-image pretraining. However, we use the pretrained convolutional part of AlexNet, as well as the pretrained acoustic network of [21] for initialisation (§IV-A2). The rest of the implementations remain the same as set out in §IV-A.

3) Few-shot evaluation: For evaluation, we only consider the few-shot classification task. We report performance over 1000 episodes, each episode consisting of five spoken query words (one for each of the few-shot classes) and five images (also one for each of the few-shot classes) in the matching set. In contrast to the English experiments in §IV and §V, each matching image can potentially belong to more than one few-shot class. The reason for this is that the few-shot classes frequently co-occur within Flickr images. To determine whether a particular image contains one of the few-shot classes, we use the ground truth text transcriptions available with the Flickr data (five text captions per image for English and one per image for Yorùbá). We mark a prediction as correct when the query audio word matches any of the words in a transcript.

TABLE VIII
THE FIVE NEW FEW-SHOT CLASSES FOR THE ENGLISH AND YORÙBÁ FLICKR EXPERIMENTS.

English	Yorùbá
boy	ọmọkùnrin
dogs	àwọn ajá
grass	koríko
rock	àpáta
water	omi

B. Experimental results

The multimodal few-shot classification results are given in Table IX. Line 2 represents the first time that multimodal few-shot word classification is performed on a real low-resource language, Yorùbá. We see that the performance of the Yorùbá model on this task (36.3%) is worse compared to the English baseline (59.4%), trained and evaluated using a similar setup.

We didn't initialise the English baseline model here by pertaining it on the background SpokenCOCO data as we did before (§IV-A) because this background data might contain paired English instances of the few-shot classes that we are using here. However, it is a fair experiment to use the SpokenCOCO English data to initialise a Yorùbá MATTNET, i.e. we use the available resources from a well-resourced language to transfer knowledge to a low-resourced setting. The results for this approach are given in line 3 of Table IX. This Yorùbá MATTNET model outperforms both the English and Yorùbá model without pretraining in terms of overall accuracy. Per-keyword classification scores are also given in Table IX. The English-pretrained Yorùbá model gives substantial improvements over the non-pretrained model across all five keyword classes, and also outperforms the English model on three of the few-shot classes.

Fig. 7 shows qualitative examples on the first episode for the English model and the Yorùbá model with pretraining. We observe that many of the predictions are accurate. Some that are deemed to be wrong could in fact be considered correct. As a reminder, our evaluation is based on the captions associated with each image, and in some cases the captions might omit a keyword even though it is actually present in the image [43]. E.g. "grass" is not mentioned in relation to the image with the boy jumping in the pool, but grass does occur in the image, and the model selects this image. So even though it is marked as a mistake in our evaluation protocol, it is actually correct. The figure also shows that the pretrained Yorùbá model seems more confident—the attention scores are higher for a given audio keyword—than the English one, showing that the pretraining

TABLE IX
FEW-SHOT CLASSIFICATION ACCURACY (%) ON THE FLICKR DATA WITH ENGLISH AND THREE YORÙBÁ VERSIONS OF MATTNET. THE YORÙBÁ PT MODEL IS A YORÙBÁ MODEL INITIALISED WITH THE WEIGHTS OF THE PRETRAINED SPOKENCOCO ENGLISH MODEL FROM §IV-A2.

Model	Mine	Average accuracy	*boy* ọmọkùnrin	*dogs* àwọn ajá	*grass* koríko	*rock* àpáta	*water* omi
Eng.	✓	59.4	56.9	28.9	64.3	78.4	68.3
Yor.	✓	36.3	18.8	23.2	35.4	49.8	54.5
Yor. pt	✓	62.0	66.9	37.4	56.6	78.0	71.3
Yor. pt	✗	29.4	47.1	30.7	29.7	21.1	18.2

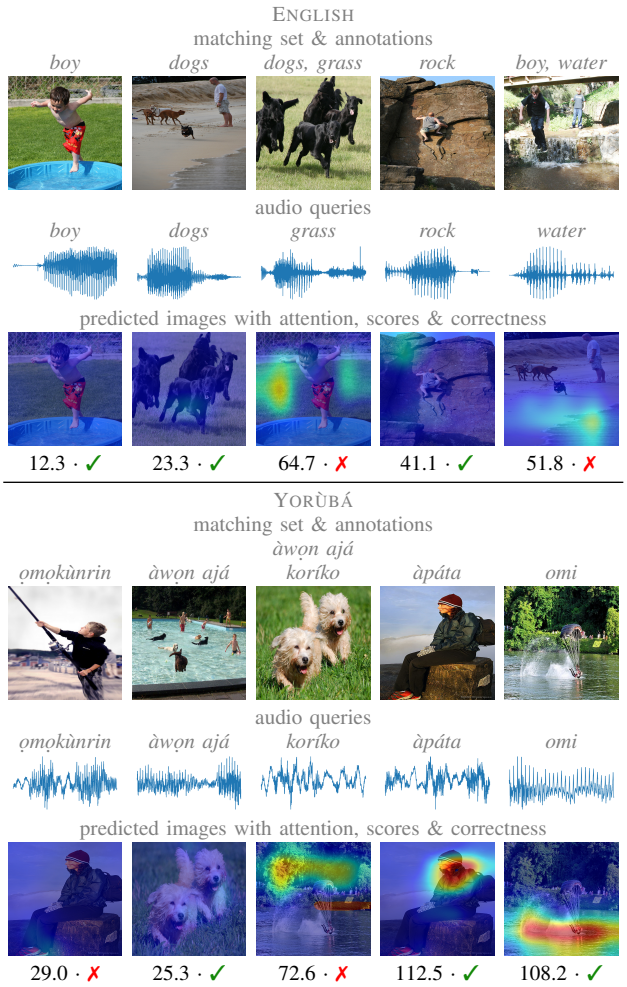

Fig. 7. Qualitative results showing the first few-shot classification episode on Flickr for English and Yorùbá. For Yorùbá we are showing results using the model with pretraining on background SpokenCOCO data. The labels are selected based on the text captions. We show attention weights overlaid on top of the images, with blue indicating the lowest scores and red the highest.

TABLE X
PRECISION SCORES (%) OF THE MINED ENGLISH AND YORÙBÁ FLICKR PAIRS, FOR IMAGE AND AUDIO MINING.

Modality	*boy* ọmọkùnrin	*dogs* àwọn ajá	*grass* koríko	*rock* àpáta	*water* omi
Images	38	46	62	19	80
English audio	95	45	87	49	97
Yorùbá audio	70	42	58	57	84

improved not only the overall performance but also the attention explanations.

Altogether, the experiments on Yorùbá shows that MATTNET can be applied to learn word-image classes from a few examples in a low-resource language when we also take advantage of a well-resourced language like English. This is similar to the findings of [44], [45].

C. Mining performance across languages

The last question that remains is: what is the effect of using an English speech system, QbERT, to mine Yorùbá audio pairs? Here we investigate the mining performance of the audio segments. In Table X we show the precision of the mined audio pairs for the two languages and five keywords. To get the precision scores, a mined audio segment is considered correct if the query keyword appears *anywhere* in the caption (that is, we do not take the temporal alignments of the mined segment into consideration).[4]

We observe that for English three keywords (*boy*, *grass* and *water*) obtain excellent results (around 90% precision), while the other two are performing more modestly (around 50% precision). To understand these results we find the most common mistakes for each of these three keywords; unsurprisingly, these mistakes are phonetically similar words: "dogs" is confused with "dog" (51 times out of the 100 mined samples) and "dock" (1), while "rock" is confused with "dog" (20) and "rocky" (13). For Yorùbá the performance across keywords has a more uniform spread and, while the top performance is lower than what we obtain on English, many of the keywords still obtain a reasonable precision.

Interestingly, the performance on the mined audio segments does not seem to necessarily correlate with the final performance. E.g. while the mined audio pairs of *rock* are among the worst, its downstream performance is the best for both the English and the Yorùbá pretrained models. A similar observation can be made for *boy*, which is accurately mined, but whose few-shot classification are poor compared to the other keywords. One reason for this might be that the accuracy of the mined image pairs for *boy* is low, as indicated in the first line of Table X which shows the precision of the mined image pairs (which is the same for both languages).

To quantify the effect of mining on classification performance, we report results for the "Yorùbá pt, no mining" model on each of the five keywords (Table IX, row 4). Comparing this to when mining is used (Table IX, row 3), we see that mining is beneficial even for keywords with modest mining performance. This suggests that more positive samples, even noisy ones, are better than fewer clean ones.

VII. CONCLUSION

Our goal was to do multimodal few-shot learning of natural images and spoken words. We proposed a novel few-shot pair mining method which we use in a new multimodal word-to-image attention model. For the scenario where the number of "shots" is small, our new model achieves higher few-shot retrieval scores than an existing model on a few-shot benchmark where an English query is used to retrieve images. We also set a competitive baseline for natural visually grounded few-shot word classification using English data.

In further analyses, we showed that a few-shot model can be used to locate occurrences of an object in an image given a spoken query, and that many of the model's mistakes are due to associating contextual information with a few-shot class, e.g. *fire-hydrant* often co-occurring with street views.

[4]For English, the differences between this metric and its stricter variant that also checks that the temporal alignments of the mined segment agree with the spoken word are minute: less than 2% absolute.

To showcase that our model can also be applied to a real low-resource language, we performed—for the first time—multimodal few-shot learning on a real low-resource language, Yorùbá. We showed that a multimodal Yorùbá few-shot model can benefit substantially from being initialised on a more substantial amount of English speech-image data.

Future work will look into extending low-resource few-shot word classification to even more classes. On both English and Yorùbá data, our analysis also revealed that the image matching step in the mining scheme might, in particular, be limiting performance. Future work should therefore use recent advances from the vision community to improve image mining.

REFERENCES

[1] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech Commun.*, 2014.

[2] I. Biederman, "Recognition-by-components: A theory of human image understanding." *Psych. Review*, 1987.

[3] G. Miller and P. Gildea, "How children learn words," *SciAM*, 1987.

[4] R. L. Gómez and L. Gerken, "Infant artificial language learning and language acquisition," *TiCS*, 2000.

[5] B. M. Lake, C.-y. Lee, J. R. Glass, and J. B. Tenenbaum, "One-shot learning of generative speech concepts," *CogSci*, 2014.

[6] O. Räsänen and H. Rasilo, "A joint model of word segmentation and meaning acquisition through cross-situational learning." *Psych. Review*, 2015.

[7] R. Eloff, H. A. Engelbrecht, and H. Kamper, "Multimodal one-shot learning of speech and images," in *Proc. ICASSP*, 2019.

[8] L. Nortje and H. Kamper, "Unsupervised vs. transfer learning for multimodal one-shot matching of speech and images," in *Proc. Interspeech*, 2020.

[9] ——, "Direct multimodal few-shot learning of speech and images," in *Proc. Interspeech*, 2021.

[10] D. Harwath, A. Recasens, D. Suris, G. Chuang, A. Torralba, and J. Glass, "Jointly discovering visual objects and spoken words from raw sensory input," in *Proc. ECCV*, 2018.

[11] H. Kamper, A. Anastassiou, and K. Livescu, "Semantic query-by-example speech search using visual grounding," in *Proc. ICASSP*, 2019.

[12] K. Olaleye and H. Kamper, "Attention-based keyword localisation in speech using visual grounding," in *Proc. Interspeech*, 2021.

[13] G. Chrupała, "Visually grounded models of spoken language: A survey of datasets, architectures and evaluation techniques," *J. Artif. Intell. Res.*, 2022.

[14] D. Merkx, S. Scholten, S. L. Frank, M. Ernestus, and O. Scharenborg, "Modelling human word learning and recognition using visually grounded speech," *Cogn. Comput.*, 2022.

[15] P. Peng, S.-W. Li, O. Räsänen, A. Mohamed, and D. Harwath, "Syllable discovery and cross-lingual generalization in a visually grounded, self-supervised speech mode," in *Proc. Interspeech*, 2023.

[16] L. Berry, Y.-J. Shih, H.-F. Wang, H.-J. Chang, H.-y. Lee, and D. Harwath, "M-SpeechCLIP: Leveraging large-scale, pre-trained models for multilingual speech to image retrieval," in *Proc. ICASSP*, 2023.

[17] T. Miller and D. Harwath, "Exploring few-shot fine-tuning strategies for models of visually grounded speech," in *Proc. Interspeech*, 2022.

[18] L. Nortje and H. Kamper, "Towards visually prompted keyword localisation for zero-resource spoken languages," in *Proc. SLT*, 2022.

[19] W.-N. Hsu, D. Harwath, C. Song, and J. Glass, "Text-free image-to-speech synthesis using learned segmental units," in *Proc ACL*, 2021.

[20] K. Olaleye, D. Oneata, and H. Kamper, "YFACC: A Yorùbá speech-image dataset for cross-lingual keyword localisation through visual grounding," in *Proc. SLT*, 2023.

[21] L. Nortje, B. van Niekerk, and H. Kamper, "Visually grounded few-shot word acquisition with fewer shots," in *Proc. Interspeech*, 2023.

[22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016.

[23] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *ACM*, 2017.

[24] H. Kamper, "Truly unsupervised acoustic word embeddings using weak top-down constraints in encoder-decoder models," in *Proc. ICASSP*, 2019.

[25] Y.-A. Chung, C.-C. Wu, C.-H. Shen, and H.-y. Lee, "Unsupervised learning of audio segment representations using sequence-to-sequence recurrent neural networks," in *Proc. Interspeech*, 2016.

[26] Y.-H. Wang, H.-y. Lee, and L.-s. Lee, "Segmental audio Word2Vec: Representing utterances as sequences of vectors with applications in spoken term detection," in *Proc. ICASSP*, 2018.

[27] N. Holzenberger, M. Du, J. Karadayi, R. Riad, and E. Dupoux, "Learning word embeddings: Unsupervised methods for fixed-size representations of variable-length speech segments," in *Proc. Interspeech*, 2018.

[28] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *ACM*, 2021.

[29] H. Kamper and B. van Niekerk, "Towards unsupervised phone and word segmentation using self-supervised vector-quantized neural networks," in *Proc. Interspeech*, 2021.

[30] S. B. Needleman and C. D. Wunsch, "A general method applicable to the search for similarities in the amino acid sequence of two proteins," *J. Mol. Biol.*, 1970.

[31] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft COCO: Common objects in context," in *Proc. ECCV*, 2014.

[32] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using Kaldi," in *Proc. Interspeech*, 2017.

[33] 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.

[34] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. ICASSP*, 2015.

[35] D. Harwath, G. Chuang, and J. Glass, "Vision as an interlingua: Learning multilingual semantic embeddings of untranscribed speech," in *Proc. ICASSP*, 2018.

[36] B. van Niekerk, L. Nortje, and H. Kamper, "Vector-quantized neural networks for acoustic unit discovery in the ZeroSpeech 2020 challenge," in *Proc. Interspeech*, 2020.

[37] D. Harwath, A. Torralba, and J. Glass, "Unsupervised learning of spoken language with visual context," in *Proc. NeurIPS*, 2016.

[38] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.

[39] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier, "Collecting Image Annotations Using Amazon's Mechanical Turk," in *NAACL HLT Workshop*, 2010.

[40] M. Hodosh, P. Young, and J. Hockenmaier, "Framing image description as a ranking task: Data, models and evaluation metrics," *JAIR*, 2013.

[41] H. Kamper, G. Shakhnarovich, and K. Livescu, "Semantic speech retrieval with a visually grounded model of untranscribed speech," *IEEE/ACM Trans. Audio Speech Lang. Process.*, 2019.

[42] D. Harwath and J. Glass, "Deep multimodal semantic embeddings for speech and images," in *Proc. ASRU*, 2015.

[43] A. C. Berg, T. L. Berg, H. Daumé, J. Dodge, A. Goyal, X. Han, A. Mensch, M. Mitchell, A. Sood, K. Stratos, and K. Yamaguchi, "Understanding and predicting importance in images," in *Proc. CVPR*, 2012.

[44] H. Kamper, Y. Matusevych, and S. Goldwater, "Improved acoustic word embeddings for zero-resource languages using multilingual transfer," *TASLP*, 2021.

[45] C. Jacobs, Y. Matusevych, and H. Kamper, "Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation," in *Proc. SLT*, 2021.