

YFACC: A YORÙBÁ SPEECH-IMAGE DATASET FOR CROSS-LINGUAL KEYWORD LOCALISATION THROUGH VISUAL GROUNDING

Kayode Olaleye¹, Dan Oneață², Herman Kamper¹

¹Stellenbosch University, South Africa

²University Politehnica of Bucharest, Romania

ABSTRACT

Visually grounded speech (VGS) models are trained on images paired with unlabelled spoken captions. Such models could be used to build speech systems in settings where it is impossible to get labelled data, e.g. for documenting unwritten languages. However, most VGS studies are in English or other high-resource languages. This paper attempts to address this shortcoming. We collect and release a new single-speaker dataset of audio captions for 6k Flickr images in Yorùbá—a real low-resource language spoken in Nigeria. We train an attention-based VGS model where images are automatically tagged with English visual labels and paired with Yorùbá utterances. This enables cross-lingual keyword localisation: a written English query is detected and located in Yorùbá speech. To quantify the effect of the smaller dataset, we compare to English systems trained on similar and more data. We hope that this new dataset will stimulate research in the use of VGS models for real low-resource languages.

Index Terms— Visually grounded speech models, keyword spotting, low-resource speech, multimodal modelling.

1. INTRODUCTION

Imagine you are a linguist tasked with documenting a low-resourced language. You collect a set of recordings from native speakers and would like to learn how common words (such as *child*, *water*, *red*) are spoken in that new language. To aid with this process, one would hope to use machine learning tools. However, developing automated models requires labelled data (e.g. audio and transcripts), which are scarce or even completely unavailable (as is the case with languages that do not have a written form). In this situation, instead, it might be reasonable to assume that we can collect images paired with their spoken description. In this paper, we investigate this scenario in which we attempt to cross-lingually localise keywords in a low-resourced language (Yorùbá). We assume that visual context is the only source of supervision.

This work is supported in part by the National Research Foundation of South Africa (grant no. 120409) and by the Romanian Ministry of Education and Research (CNCS-UEFISCDI, number PN-III-P1-1.1-PD-2019-0918).

Models that learn from images and their spoken captions are known as visually grounded speech (VGS) models [1–11]. While previous work on VGS has shown promise for performing a number of downstream tasks in low-resource conditions, very few papers validate the approach in a realistic setting—most papers use unlabelled English speech or data from another high-resource language [7, 12] to simulate a low-resource setting.

To address this shortcoming in VGS research, we introduce a new speech-image dataset in a low-resourced language, Yorùbá. The Yorùbá language is one of the three official languages in Nigeria, and even though it is spoken natively by more than 44M speakers, there are only a handful of high-quality datasets available in this language [13–16]. Our dataset is built by augmenting the Flickr8k dataset [17, 18] with spoken captions: from a single Yorùbá speaker, we record 6k Yorùbá utterances corresponding to translations of the original English image captions. We also provide manual alignments

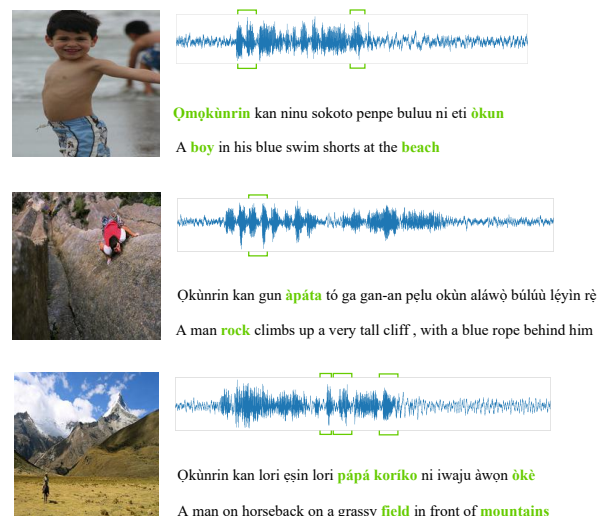


Fig. 1. Samples from the new Yorùbá Flickr Audio Caption Corpus (YFACC): images from the Flickr8k dataset (left) are paired with Yorùbá audio captions (right), which are based on the original English captions. For a subset of keywords (green) we provide temporal alignments for evaluation.

with start and end timestamps for 67 selected keywords in each of the 500 test utterance. We call the resulting dataset the Yorùbá Flickr Audio Captions Corpus (YFACC). Examples of samples from YFACC are shown in Figure 1. A major advantage of extending Flickr8k is that it allows reusing existing annotations (English recordings [3], translations in other languages [19, 20]), which could be used to construct richer and potentially more interesting tasks in future work.

In this paper we specifically use YFACC for the tasks of cross-lingual keyword detection and localisation: given an English text query, we detect whether the query occurs in Yorùbá speech, and if it is detected, we localise where in the utterance the query occurs. This is a novel scenario that resembles language documentation. Prior work has considered cross-lingual keyword detection [21] (whether the query appears *anywhere* in an utterance), but localisation was not attempted (and they considered the high-resource English–German language pair). Other work has considered VGS-based localisation [22, 23], but in these monolingual studies the query keywords came from the *same* language as the search utterances. To our knowledge, we are the first to attempt cross-lingual keyword localisation; while more challenging, this task is also potentially more impactful in developing applications for linguists. Our approach applies the model of [23] to the cross-lingual setting: images are automatically tagged with English visual labels, which serve as targets for an attention-based model that takes Yorùbá speech as input. This enables the Yorùbá speech model to be queried with English text.

Apart from the inherent challenges in the cross-lingual retrieval task, the English–Yorùbá pair is also a challenging combination, as the two languages belong to different families (Indo-European and Nigerian-Congo) and they are associated with very different cultures. Through systematic experimentation, we quantify the impact of this language mismatch: we analyse how well the selected keywords match across languages, we compare to a monolingual English VGS model trained on similar amounts of data to YFACC, and we transfer English speech representations to Yorùbá (resulting in improved retrieval performance).

To summarise, our contributions include: (i) the release of a new speech–image dataset in Yorùbá, (ii) a baseline system for cross-lingual keyword detection and localisation, and (iii) a multi-facet cross-lingual analysis of systems. We hope that YFACC will stimulate research in the use of VGS models for solving speech tasks in real low-resource languages. We release the data at <https://www.kamperh.com/yfacc/>.

2. RELATED WORK

Visually-grounded speech models. There is a growing body of work considering how speech processing systems can be developed using images paired with spoken captions [2–4, 6, 7, 9, 10, 21, 24]. One line of work trains a model to project images and speech into a joint embedding space [4, 25]. Another line

uses an external image tagger with a fixed vocabulary to obtain soft text labels, which are then used as targets for a speech network that maps speech to keyword labels [26, 27]. Our work falls in the second category and investigates how these systems perform in a real low-resource environment.

Cross-lingual speech-keyword retrieval. Early work in cross-lingual retrieval adopted a direct approach of using text transcriptions for training the system instead of visual grounding. E.g. [28] proposed to cascade automatic speech recognition with text-based cross-lingual information retrieval [29]. However, this approach is only possible when transcribed speech is available in the target language to build a recogniser. Recent work has proposed models that can translate speech in one language directly to text in another [30–33], but these methods also require transcribed data: parallel speech with translated text. [21] attempted cross-lingual keyword spotting without the need of transcriptions by using the visual information as the link between the languages. Our approach shares this core idea, but differs in terms of the languages used and the task tackled (as motivated in Sec. 1).

Yorùbá speech corpora. Up until very recently, there were only two open-source Yorùbá speech datasets: the Lagos-NWU Yorùbá Speech Corpus [13], consisting of approximately 2.75 hours recorded by 17 male and 16 female speakers at 16 kHz, and the corpus of Gutkin *et al.* [14], consisting of roughly 4 hours of 48 kHz recordings from 36 male and female volunteers. Concurrent work [15, 16], acknowledging the importance of creating datasets for African languages, has released more sizeable datasets geared towards building text-to-speech systems; Table 1 gives their statistics. We believe that YFACC, being a multimodal speech-image dataset, will make a further useful contribution to this short list.

3. THE YORÙBÁ FLICKR AUDIO CAPTION CORPUS (YFACC)

We introduce the new Yorùbá Flickr Audio Caption Corpus (YFACC). The YFACC dataset extends the Flickr8k image–text dataset [17, 18] to Yorùbá with three modalities: Yorùbá translations of 6k of its captions; corresponding spoken recordings of these translations; temporal alignments of 67 Yorùbá keywords for a subset of 500 of the captions.

3.1. Data collection

We selected a random set of 6k images from the Flickr8k dataset and for each image we picked one out of the five associated captions. The resulting English captions were manually translated to Yorùbá by two native speakers. Partial diacritics were added to aid the recording process. All translations were recorded from a single male native speaker. A primary motivation for collecting speech from only a single speaker is that this is representative of a real language documentation setting, where a linguist will only have access to a small number of

Table 1. A comparison of the open-source Yorùbá speech corpora in terms of the available modalities and quantitative characteristics. The new YFACC datasets provides alignments for a subset of 67 keywords.

Dataset	Modalities					Speakers number	Duration hours	Utterances number	Sampling kHz
	Speech	Text (yo)	Text (en)	Images	Alignments				
van Niekerk et al. [13]	✓	✓				33	2.75	4316	16
Gutkin et al. [14]	✓	✓				36	4.00	3583	48
Meyer et al. [15]	✓	✓					33.30	10228	48
Ogayo et al. [16]	✓	✓					18.04	10978	
YFACC (ours)	✓	✓	✓	✓	partial	1	13.3	6000	48

speakers. Recordings were made with a BOYA BY-M1DM dual omni-directional lavalier microphone using a sampling rate of 48 kHz and two channels.

The 6k utterances were split into train (5k), development (500) and test (500) sets. For the 500 test utterance, a native speaker produced temporal alignments for a set of 67 Yorùbá keywords corresponding to the 67 English keywords defined in [26]. These keywords were selected such that they are visually groundable and provide high inter-annotator agreement, so they are suitable to evaluate the localisation algorithm. We prepare the alignments with Praat [34]. Alignments include the start time, end time, and duration of each of the 67 keywords. Samples from YFACC are shown in Figure 1.

3.2. Relation to Flickr8k and other Yorùbá datasets

Our dataset is based on the Flickr8k dataset [17, 18], which is a bimodal (image and text) dataset consisting of 8k images and 40k captions (each image is annotated with a textual description by five different people). The images (and consequently their captions) generally depict relations and actions involving people or animals. Flickr8k was initially used for image captioning and cross-modal retrieval, but was subsequently extended with translations [19, 20], spoken captions [3], and bounding box annotations for the visual concepts [35]. Of particular interest to us is the spoken version, Flickr Audio Caption Corpus (FACC), introduced by Harwath and Glass [3]. FACC was obtained using Amazon’s Mechanical Turk to crowdsource English audio captions for each of the text captions. The FACC dataset subsequently enabled visually-grounded tasks, such as image–speech alignment [3], semantic speech retrieval [26], image-to-speech synthesis [36] and speech-to-image generation [37].

Compared to previous Yorùbá datasets, ours is different in that it has a richer set of modalities, as shown in Table 1. By extending Flickr8k we are able to reuse its existing annotations and build a multi-modal multi-lingual dataset, which, similarly to FACC, allows for a wide variety of tasks to be defined. In this current paper we specifically consider the task of visually-grounded cross-lingual keyword localisation. Example tasks that can be considered in the future include speech-to-speech machine translation or grounding visual concept in images

from Yorùbá speech. The fact that the recordings are done by a single speaker and at high quality (48 kHz), also enables the future development of text-to-speech technology for Yorùbá.

3.3. Cross-linguistic differences

Yorùbá is a very different language from English both linguistically and culturally. In what follows we discuss some of the specific aspects as they relate to YFACC.

Tone marks. In Yorùbá, pronunciation and writing involves using three types of tone marks [38]: *dò*, a falling tone (grave accent); *re*, a flat tone (no accent); *mí*, a rising tone (acute accent). These marks are applied to the top of the seven Yorùbá vowels (*a, e, ẹ, i, o, ọ, u*), and the four nasal vowels (*ẹn, in, ọn, un*). The tone marks are key to understanding Yorùbá, since without them many words become ambiguous. E.g. *okun* can stand for *okun* (beach), *okùn* (rope) or *okun* (strength), and *ọkọ* can stand for *ọkọ* (car), *ọkọ* (hoe) or *ọkọ* (husband). In practice including all the marks when writing is a time consuming activity. So for our dataset we provide diacritics only for the 67 keywords, since only these are used for evaluation. However, in the future it might be possible to generate the marks for the remaining words using an automated approach [39].

Keywords across the two languages. Even if YFACC is built by translating the Flickr8k dataset, a keyword that appears in the English caption is not guaranteed to appear in the corresponding Yorùbá translation, and vice versa. For example, *A blonde woman appears to wait for a ride* is translated to Yorùbá as *Obìnrin bílòndì ndúró de ọkọ* (which back-translates to *A blonde woman stands for a car*), making the keywords *stands* and *car* appear in the Yorùbá caption. We quantify the common appearances of the keywords across the two languages in Figure 2. Intuitively, if an English and an Yorùbá word tend to co-occur (i.e., appear in the corresponding English–Yorùbá translation samples, as is the case with the *ride–ọkọ* pair in the previous example), that pair will yield a large score. Indeed, we observe a strong correlation along the diagonal, suggesting that English keywords do correlate with their translations, but there are also some exceptions. Some Yorùbá words have multiple meanings, so they co-occur with more than a single English word, e.g. *òkun*, which means both *beach* and *ocean*. Conversely, some English words have multiple meanings: *foot-*

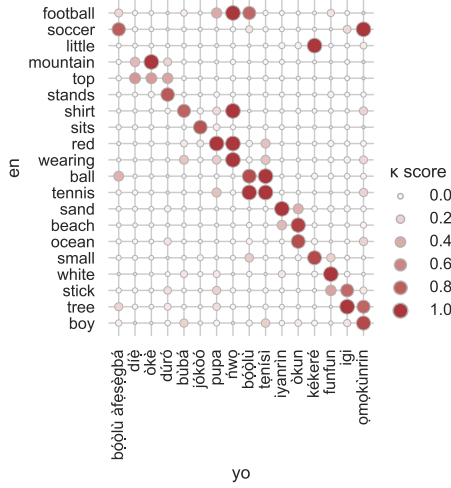


Fig. 2. Co-occurrences of English–Yorùbá keyword pairs on the 500 test utterances in YFACC computed as normalised Cohen’s κ score. For readability, we pick keywords from the 67-word vocabulary.

ball is not just the sport, but also the *ball*; *little* refers not only to size, but also to age (*little boy*; *omọ̀kùnrin kẹ́kẹ́rẹ̀*). Finally, some correlations appear just because words co-occur together in both languages: *nwo* (*wearing*) correlates with *shirt*, *red*, and *football*.

Cultural differences. Starting from an English dataset may induce a mismatch regarding the cultural aspects, as discussed in [40]. Most of the 67 selected keywords in Flickr8k pertain to general cross-cultural concepts (such as *children*, *young*, *swimming*), but there are some of the English words that have no indigenous Yorùbá translation. E.g. words like *jacket* or *tennis* are usually translated based on their English pronunciation by adding extra marks to the English word: *jákẹ̀tì* and *tẹ́nìsì*. Perhaps the only keyword that presents little interest to the typical Nigerian is the English word *skateboard*.

4. CROSS-LINGUAL KEYWORD LOCALISATION WITH A VISUALLY-GROUNDED SPEECH MODEL

We specifically consider the task of cross-lingual keyword localisation: given an English keyword and a spoken utterance in Yorùbá, we want to automatically find the location in the audio of the Yorùbá word that corresponds to the given English keyword. If the keyword does not occur, the model should indicate that the keyword was not detected. Crucially, we assume that the only source of supervision comes from a visual channel. Concretely in our case we assume we have access to a set of images associated with corresponding audio captions (YFACC). In order to be able use this weak source of supervision, we follow an almost identical approach to the monolingual VGS-based localisation method of [23], which

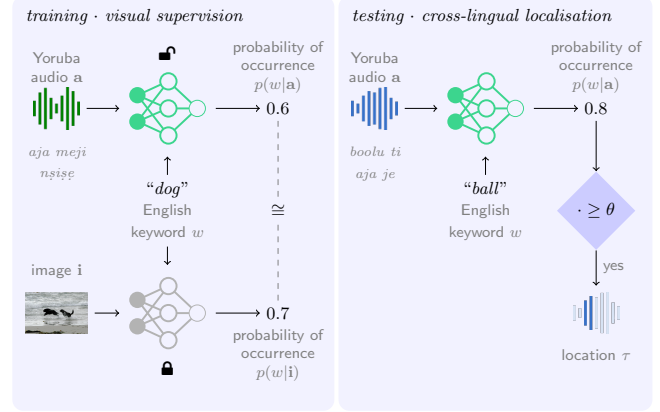


Fig. 3. Cross-lingual keyword localisation with a visually-grounded speech model. *Left:* At train time, the speech model is trained to predict whether an English keyword occurs in a Yorùbá audio utterance based on the visual supervision of a pretrained and frozen (English) image model. *Right:* At test time, given an audio utterance in Yorùbá, the model outputs the probability of an English keyword occurring in the utterance. If this probability is over a predefined threshold θ , then the model additionally locates the keyword in the utterance based on the attention weights in the network.

we adapt to the cross-lingual setting.

4.1. Training: Visual supervision

Our model consists of an audio network that processes Yorùbá speech and outputs the probability that a given English keyword occurs (anywhere) in the utterance. This audio model is trained by cross-modal supervision: it learns to reproduce the predictions of a pretrained image network on the associated image for a predefined set of visual categories. The weights of the audio network are optimised through the binary cross-entropy loss between the image and audio predictions; the loss is averaged over the keywords in the vocabulary. The categories are defined on English keywords (such as *dog*, *child*, *ball*), but they also correspond to Yorùbá words which occur in the audio, with the visual channel linking the two languages. See Figure 3 (left) for an illustration of the training process.

4.2. Testing: Cross-lingual localisation

At test time, given a Yorùbá utterance—no images are needed for inference—our model outputs the probability of occurrence of a given English keyword from the visual vocabulary. When this probability is over a certain threshold θ , we assume that the keyword occurs in the utterance. Once a keyword is detected, we also want to locate the corresponding Yorùbá word. To obtain the time location of the Yoruba word, we leverage the extra structure of our model, namely its attention weights over the input utterance. We pick as the most probable location of

the input keyword the time stamp with the highest weight. See Figure 3 (right) for an illustration of the testing phase.

4.3. Architecture: Attention network

The audio network consists of four subcomponents: an audio encoder, a keyword encoder, an attention layer and a classification network. The audio encoder—a convolutional network in our case—embeds the Yorùbá audio to a sequence of feature vectors, while the keyword encoder assigns an embedding vector to the English keyword using a learnable lookup table. The audio features are pooled temporally into a single vector based on the keyword embedding using the attention layer. Finally, the aggregated feature vector is then projected to a classification score using a small multi-layer perceptron network. As mentioned in the previous subsection, we use the location of the maximum attention scores as the location prediction of the query keyword. This architecture is based on the work of [23], which in turn draws inspiration from [41], an attention-based graph-convolutional network for visual keyword spotting in sign language videos.

5. EXPERIMENTAL SETUP

5.1. Datasets

In our experiments, apart from the new YFACC dataset (Sec. 3) we also use the previous English FACC dataset (Sec. 3.2) [3] for the monolingual English VGS systems which we compare to. We experiment with two variants of FACC that differ in the amount of training data: *en-30k*, which contains all the original 30k training samples, and *en-5k*, which contains a subset of 5k training samples to match the size of YFACC. The test data used for models trained on FACC are the 500 English utterances corresponding to the YFACC test set.

5.2. Model architecture and implementation details

We use the same architecture as in [23], but carry out a hyperparameter search over the learning rate, activations functions and size of last layer. This search is done based on F1 keyword detection score on the development set of FACC. Compared to the original architecture in [23], the only change is increasing the last layer’s size from 4096 to 8192; this improved development-set F1 from 14.4% to 19.4%.

We downsample YFACC to 16 kHz to match FACC. All utterances are encoded as mel-frequency cepstral coefficients, which are further augmented at train time using SpecAugment [42]. The model’s visual targets are computed with a VGG-16 network [43] trained on MSCOCO [44] for 67 keywords as described in [26]. Our VGS models are optimised with Adam [45] using a learning rate of 10^{-4} . The PyTorch implementation is publicly available.¹

¹https://github.com/kayodeolaleye/keyword_localisation_speech

5.3. Evaluation of cross-lingual keyword localisation

Our main metric is precision, which is computed as the ratio of true positive samples to retrieved samples. A sample is *retrieved* if the detection score for the given English keyword is greater than a threshold θ . A sample is a *true positive* if (1) it is retrieved and (2) the given English keyword is correctly localised in the Yorùbá utterance, i.e., the predicted location falls within the interval of the corresponding spoken Yorùbá keyword. We refer to this metric as *actual keyword localisation* [23]. For all experiments, we set θ to 0.5, which gave best performance on the FACC development set.

6. EXPERIMENTAL RESULTS AND ANALYSIS

Our goal is to investigate to what extent cross-lingual keyword localisation is possible with a VGS model trained in an extreme low-resource setting. We specifically consider the case where we have only 13.3 hours of untranscribed Yorùbá speech with corresponding images.

English–Yorùbá cross-lingual results are given at the top of Table 2. The cross-lingual model *yo-5k* obtains a performance of about 16.0% for our main task, actual keyword localisation (Sec. 5.3). Although this appears modest when viewed in isolation, it is much better than the performance of a random model (0.1%). More importantly, the result falls within the range of results ($21.3 \pm 4.4\%$) that are obtained in a monolingual setting (querying English keywords in English speech) with a model trained on equal number of English utterances (*en-5k*). It can be argued that the cross-lingual setting is more challenging because the visual keywords are more English-centric and do not account for the cultural and linguistic differences to Yorùbá (Sec. 3.3). On the other hand, YFACC is easier in the sense its a single-speaker dataset, whereas the FACC dataset involves multiple speakers. Next we consider aspects that can help improve performance.

Amount of training data. While for the cross-lingual setting we are limited to the 5k training utterances available in YFACC, we can extrapolate the performance on more data based on corresponding monolingual setting using FACC. In Table 2 we see that training with six times more data, unsurprisingly, yields considerable improvements, from 21.3% (*en-5k*) to 26.9% (*en-30k*). However, while increasing the training size is the most straightforward way to improve performance, it might also be the most challenging in practice.

Speech representations. Training a model from scratch in low-resource conditions is challenging. A common practice is to start from a pretrained model and to then fine-tune it on the downstream task. Here we do this by initialising the training on Yorùbá from a VGS model trained on all the image–English audio pairs in FACC. Such English pairs are arguably more readily available than Yorùbá utterances. We observe a strong boost in performance—from 16.0% (*yo-5k*) to 22.8% (*yo-5k-init*).

Table 2. Results of visually-grounded speech models on cross-lingual tasks (querying English in Yorùbá audio) on the proposed YFACC dataset (first three rows) and on monolingual tasks (querying English in English audio) on the FACC dataset (last two rows). The models are evaluated for the main task of keyword localisation (actual localisation) and two other easier tasks (oracle localisation and keyword detection), which provide upper bounds on the actual localisation performance. The experiments are repeated three times by changing the random seed and we report the mean performance and its standard deviation.

	Keyword language	Audio language	Training num. utts.	Initialisation	Actual localisation precision	Oracle localisation accuracy	Keyword detection precision
random	Random baseline				0.1 ± 0.0	5.3 ± 0.9	2.4 ± 0.2
yo-5k	English	→ Yorùbá	5k	random	16.0 ± 1.6	23.1 ± 3.0	26.7 ± 6.4
yo-5k-init	English	→ Yorùbá	5k	en-30k	22.8 ± 4.3	35.0 ± 2.5	33.3 ± 7.1
en-5k	English	→ English	5k	random	21.3 ± 4.4	27.3 ± 2.8	29.5 ± 5.0
en-30k	English	→ English	30k	random	26.9 ± 1.8	44.0 ± 0.9	39.2 ± 1.7

Decoupling the tasks. Keyword localisation involves two steps: (1) detecting whether the query keyword appears (somewhere) in the utterance, and, if so, (2) localising the keyword in the given utterance. In order to understand which component of our approach requires improvements, it is instructive to independently evaluate each of these steps. To do that, we assume that one of the steps is perfect, resulting in two new tasks: oracle localisation (assumes that detection is perfect) and keyword detection (assumes that localisation is perfect). Whilst not directly comparable, the results for the two tasks (last two columns in Table 2) indicate that both components can be further improved. Also note that the previous conclusions in terms of relative performance between the models still hold on these new tasks.

Individual keyword performance. The performance varies widely across the 67 keywords in the vocabulary. We observe that there are keywords with good performance (better than the average), such as *brown* (*búrúún*; 100%), *bike* (*kèké*; 94.1%) or *grass* (*koríko* 90.9%). But there are many others on which the model struggles. The reason for the latter include poor visual grounding (e.g., *camera*, *wearing*, *large*) and confusion between semantically related concepts (*riding* often retrieves *kèké*, i.e. *bicycle*; *swimming* retrieves *odò adágún*, i.e. *pool*). Paying close attention to the particularities of the individual keywords can serve as inspiration for future improvements.

Qualitative results. To better understand the failure modes of the cross-lingual model, we qualitatively assess the results. Figure 4(a) and (b) show correct localisations of English keywords in Yorùbá utterances: the model fires within the frames containing the spoken word *òpópónà* when prompted with the English keyword *street* (4a) and fires within the frames of *òkun* when prompted with the keyword *ocean*. Figures 4(c) and (d) show failure cases. Figure 4(c) shows that the model also sometimes localises a semantically related word in the utterance: the spoken Yorùbá word *òkun* (*ocean*) is assigned the highest localisation score when prompted with the English keyword *water*. Figure 4(d) is an outright failure.

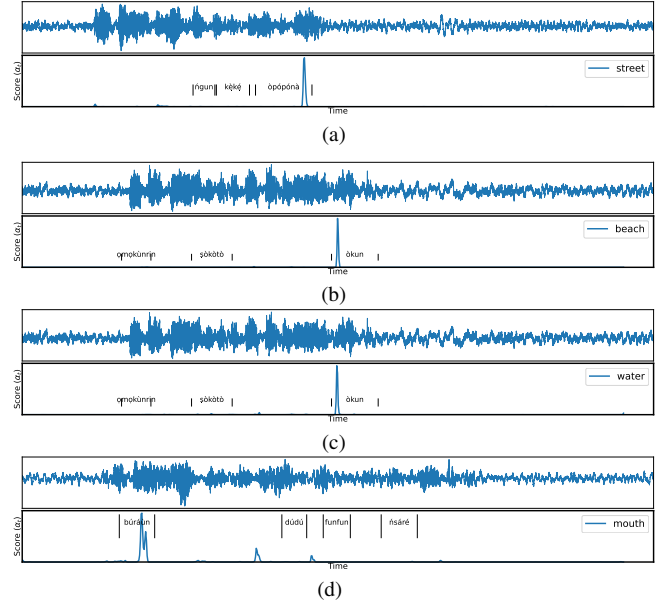


Fig. 4. Examples of cross-lingual localisation with the yo-5k VGS model. The English query keyword is shown on the right.

7. CONCLUSION

We introduced a new single-speaker dataset of Yorùbá spoken captions for Flickr images. To showcase what is possible with such a small dataset (13.3 hours, 6k speech-image pairs) in a real low-resource setting, we used the data for cross-lingual keyword localisation using a visually grounded speech model, trained without any transcribed audio. By initialising the model with an English model trained on more data, we were also able to obtain some improvements in keyword localisation and detection performance. This dataset with the accompanying baselines will allow researchers to evaluate their visually grounded speech models in a real low-resource setting.

8. REFERENCES

- [1] J. Driesen, H. Van hamme, and W. B. Kleijn, “Learning from images and speech with non-negative matrix factorization enhanced by input space scaling,” in *Proc. SLT*, 2010.
- [2] G. Synnaeve, M. Versteegh, and E. Dupoux, “Learning words from images and speech,” *NeurIPS Workshop Learn*, 2014.
- [3] D. Harwath and J. Glass, “Deep multimodal semantic embeddings for speech and images,” in *Proc. ASRU*, 2015.
- [4] D. Harwath, A. Torralba, and J. Glass, “Unsupervised learning of spoken language with visual context,” in *Proc. NeurIPS*, 2016.
- [5] D. Harwath and J. Glass, “Learning word-like units from joint audio-visual analysis,” in *Proc. ACL*, 2017.
- [6] D. Harwath, A. Recasens, D. Surís, G. Chuang, A. Torralba, and J. Glass, “Jointly discovering visual objects and spoken words from raw sensory input,” in *Proc. ECCV*, 2018.
- [7] D. Harwath, G. Chuang, and J. Glass, “Vision as an interlingua: Learning multilingual semantic embeddings of untranscribed speech,” in *Proc. ICASSP*, 2018.
- [8] R. Eloff, H. A. Engelbrecht, and H. Kamper, “Multi-modal one-shot learning of speech and images,” in *Proc. ICASSP*, 2019.
- [9] D. Harwath and J. Glass, “Towards visually grounded sub-word speech unit discovery,” in *Proc. ICASSP*, 2019.
- [10] D. Harwath, W.-N. Hsu, and J. Glass, “Learning hierarchical discrete linguistic units from visually-grounded speech,” in *Proc. ICLR*, 2019.
- [11] L. Nortje and H. Kamper, “Unsupervised vs. transfer learning for multimodal one-shot matching of speech and images,” in *Proc. Interspeech*, 2020.
- [12] Y. Ohishi, A. Kimura, T. Kawanishi, K. Kashino, D. Harwath, and J. Glass, “Trilingual semantic embeddings of visually grounded speech with self-attention mechanisms,” in *Proc. ICASSP*, 2020.
- [13] D. van Niekerk, E. Barnard, O. Giwa, and A. Sosimi, “Lagos-NWU Yoruba speech corpus,” <https://repo.sadilar.org/handle/20.500.12185/431>, 2015.
- [14] A. Gutkin, I. Demirsahin, O. Kjartansson, C. Rivera, and K. Túbosún, “Developing an open-source corpus of Yoruba speech,” in *Interspeech*, 2020.
- [15] J. Meyer, D. I. Adelani, E. Casanova, A. Öktem, D. W. J. Weber, S. Kabongo, E. Salesky, I. Orife, C. Leong, P. Ogayo *et al.*, “BibleTTS: a large, high-fidelity, multilingual, and uniquely African speech corpus,” in *Proc. Interspeech*, 2022.
- [16] P. Ogayo, G. Neubig, and A. W. Black, “Building African voices,” *arXiv preprint arXiv:2207.00688*, 2022.
- [17] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier, “Collecting image annotations using Amazon’s Mechanical Turk,” *NAACL HLT Workshop*, 2010.
- [18] M. Hodosh, P. Young, and J. Hockenmaier, “Framing image description as a ranking task: Data, models and evaluation metrics (extended abstract),” in *Proc. IJCAI*, 2015.
- [19] D. Elliott, S. Frank, K. Sima’an, and L. Specia, “Multi30k: Multilingual English-German image descriptions,” *ACL Vis. and Lang. Workshop*, 2016.
- [20] X. Li, W. Lan, J. Dong, and H. Liu, “Adding Chinese captions to images,” in *Proc. ICMR*, 2016.
- [21] H. Kamper and M. Roth, “Visually grounded cross-lingual keyword spotting in speech,” in *Proc. SLTU*, 2018.
- [22] K. Olaleye and H. Kamper, “Attention-based keyword localisation in speech using visual grounding,” in *Proc. Interspeech*, 2021.
- [23] K. Olaleye, D. Oneata, and H. Kamper, “Keyword localisation in untranscribed speech using visually grounded speech models,” *IEEE JSTSP*, 2022.
- [24] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” in *Proc. Interspeech*, 2017.
- [25] G. Chrupała, L. Gelderloos, and A. Alishahi, “Representations of language in a model of visually grounded speech signal,” in *Proc. ACL*, 2017.
- [26] H. Kamper, G. Shakhnarovich, and K. Livescu, “Semantic speech retrieval with a visually grounded model of untranscribed speech,” *IEEE/ACM Trans. Acoust., Speech, Signal Process.*, 2019.
- [27] A. Pasad, B. Shi, H. Kamper, and K. Livescu, “On the contributions of visual and textual supervision in low-resource semantic speech retrieval,” in *Proc. Interspeech*, 2019.
- [28] P. Sheridan, M. Wechsler, and P. Schäuble, “Cross-language speech retrieval: Establishing a baseline performance,” in *Proc. SIGIR*, 1997.

- [29] D. W. Oard and A. R. Diekema, “Cross-language information retrieval.” *ARIST*, 1998.
- [30] L. Duong, A. Anastasopoulos, D. Chiang, S. Bird, and T. Cohn, “An attentional model for speech translation without transcription,” in *Proc. NAACL*, 2016.
- [31] S. Bansal, H. Kamper, A. Lopez, and S. Goldwater, “Towards speech-to-text translation without speech recognition,” in *Proc. EACL*, 2017.
- [32] 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.
- [33] A. Bérard, L. Besacier, A. C. Kocabiyikoglu, and O. Pietquin, “End-to-end automatic speech translation of audiobooks,” in *Proc. ICASSP*, 2018.
- [34] P. Boersma and V. Van Heuven, “Speak and unspeak with PRAAT,” *Glott Int.*, 2001.
- [35] B. A. Plummer, L. Wang, C. M. Cervantes, J. C. Caicedo, J. Hockenmaier, and S. Lazebnik, “Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models,” *IJCV*, 2015.
- [36] M. Hasegawa-Johnson, A. Black, L. Ondel, O. Scharenborg, and F. Ciannella, “Image2speech: Automatically generating audio descriptions of images,” in *Proc. IC-NLSP*, 2017.
- [37] J. Li, X. Zhang, C. Jia, J. Xu, L. Zhang, Y. Wang, S. Ma, and W. Gao, “Direct speech-to-image translation,” *IEEE JSTSP*, 2020.
- [38] J. B. Wood, *Notes on the construction of the Yoruba language*. University of Exeter, 1879.
- [39] I. Orife, D. I. Adelani, T. E. Fasubaa, V. Williamson, W. F. Oyewusi, O. Wahab, and K. Tubosun, “Improving Yoruba diacritic restoration,” *arXiv preprint arXiv:2003.10564*, 2020.
- [40] D. Hershcovich, S. Frank, H. C. Lent, M. de Lhoneux, M. Abdou, S. Brandl, E. Bugliarello, L. C. Piqueras, I. Chalkidis, R. Cui, C. Fierro, K. Margatina, P. Rust, and A. Sogaard, “Challenges and strategies in cross-cultural NLP,” in *Proc. ACL*, 2022.
- [41] N. C. Tamer and M. Saraçlar, “Keyword search for sign language,” in *Proc. ICASSP*, 2020.
- [42] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. Interspeech*, 2019.
- [43] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [44] 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.
- [45] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.