

# TOWARDS VISUALLY PROMPTED KEYWORD LOCALISATION FOR ZERO-RESOURCE SPOKEN LANGUAGES

Leanne Nortje and Herman Kamper

MediaLab, Electrical & Electronic Engineering, Stellenbosch University, South Africa

## ABSTRACT

Imagine being able to show a system a visual depiction of a keyword and finding spoken utterances that contain this keyword from a zero-resource speech corpus. We formalise this task and call it visually prompted keyword localisation (VPKL): given an image of a keyword, detect and predict where in an utterance the keyword occurs. To do VPKL, we propose a speech-vision model with a novel localising attention mechanism which we train with a new keyword sampling scheme. We show that these innovations give improvements in VPKL over an existing speech-vision model. We also compare to a visual bag-of-words (BoW) model where images are automatically tagged with visual labels and paired with unlabelled speech. Although this visual BoW can be queried directly with a written keyword (while our’s takes image queries), our new model still outperforms the visual BoW in both detection and localisation, giving a 16% relative improvement in localisation F1.

**Index Terms**— Visually grounded speech models, keyword localisation, speech-image retrieval.

## 1. INTRODUCTION

How can we search a speech collection in a zero-resource language where it is impossible to obtain text transcriptions (e.g. unwritten languages)? One way in which recent research is addressing this problem is to use vision as a weak form of supervision: speech systems are built on images paired with unlabelled spoken captions—removing the need for text [1].

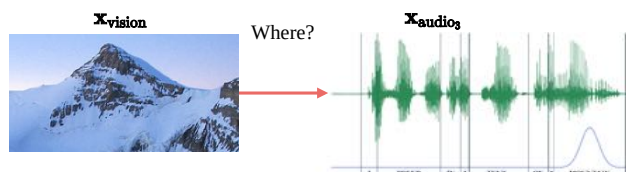
In this paper we specifically introduce the new visually prompted keyword localisation (VPKL) task: a model is given an image depicting a keyword—serving as an image query—and is prompted to detect whether the query occurs in a spoken utterance. If the keyword is detected, the model should also determine where in the utterance the keyword occurs. E.g. the model is shown an image of a *mountain* and asked whether it occurs in the spoken caption: “a hiker in a tent on a mountain”. The model should also say where in the utterance *mountain* occurs (if it is detected), as shown in Fig. 1. To do this, we need a multimodal model that can compare images or image regions to spoken utterances.

Leanne Nortje is funded through a DeepMind PhD scholarship.

In the last few years a range of speech-vision models have been proposed [1–8]. Most were developed for retrieving whole images given a whole spoken caption as query (or vice versa). Image-caption retrieval is different from VPKL—in the latter, the query is typically a depiction of an isolated object or concept and we want to detect and localise this query *within* an utterance (rather than retrieving a whole spoken caption). Nevertheless, with slight modification, we can use an image-caption retrieval model for VPKL. We show that this performs poorly, presumably because of the mismatch between the training objective and the test-time VPKL task.

As a result, we propose a novel localising attention mechanism and a new keyword sampling scheme. First, for the attention mechanism, we combine the idea of *matchmaps* [2] with a more explicit form of within-utterance attention [9, 10]. Second, for the sampling scheme, we can use a visual tagger to automatically tag training images with text labels of words likely occurring in the image. From these generated tags, we can sample positive and negative image-caption pairs which contain the same or different keywords. E.g. while originally we could have a spoken caption “hikers going up a mountain slope” paired only with a single image, we could now also pair this utterance with the spoken caption “a boy and his dad on a mountain.” This would encourage the model to not only focus on utterances as a whole, but also learn within-utterance distinctions between keywords. Note that in this paper we mainly consider an idealised case using the captions’ text transcriptions to sample positive and negative pairs (simulating an ideal visual tagger).

In this setting, we show that both innovations lead to improvements in VPKL over an image-caption retrieval model. We also compare to a visual bag-of-words (BoW) model [9, 11, 12] which is queried with written keywords in-




**Fig. 1.** The goal in visually prompted keyword localisation is to locate a given query keyword (given as an image) within a spoken utterance.

stead of images. This model is trained using a visual tagger to generate textual BoW labels for training images. These labels are then used to train a keyword detection model [12]. While a written keyword arguably gives a stronger query signal than an image, we show that when combining our new attention mechanism with our new sampling scheme, we outperform the visual BoW. Further analysis shows that the distribution of keywords that our model is able to localise is much smaller than that of the visual BoW model. We attribute this to image queries sometimes depicting more than one keyword. Through further analyses, we also show that the model’s performance decreases when tasked with learning a larger set of keywords. We also present initial experiments where a real image tagger is used to produce positive and negative examples for our training scheme—highlighting additional challenges for future work.

## 2. NEW TASK: VISUALLY PROMPTED KEYWORD LOCALISATION

The approach of directly training on image-speech pairs is motivated by children having access to image and speech signals when acquiring their native language [13–19]. To learn words, they can use the co-occurrences of spoken words with visual objects, and vice versa [20]. Eventually humans can establish if and where a word depicted by its visual representation, is uttered—without ever requiring transcriptions. Drawing inspiration from humans, we introduce the new task of visually prompted keyword localisation (VPKL). This task is very similar to the task of textual keyword detection, where a model is given a written query keyword and asked to detect (and possibly locate) occurrences of the keyword in a search collection [9, 11, 12, 21–25]. Instead of a written keyword, in VPKL the query is an image of an object or concept.

Formally, VPKL involves both detection and localisation. Detection is illustrated in Fig. 2. A model is given an image query  $\mathbf{x}_{\text{vision}}$ , which depicts a keyword, and asked whether the keyword occurs in an utterance  $\mathbf{x}_{\text{audio}}$ . For localisation, if the model detects the image query  $\mathbf{x}_{\text{vision}}$  in  $\mathbf{x}_{\text{audio}}$ , the model is prompted to identify where in  $\mathbf{x}_{\text{audio}}$  the keyword occurs. E.g. in Fig. 1 the model is asked whether the *mountain* in  $\mathbf{x}_{\text{vision}}$ , occurs in  $\mathbf{x}_{\text{audio}_3}$ . After the model detected the key-



**Fig. 2.** The goal in visually prompted detection is to detect whether a given query keyword (given as an image) occurs anywhere within a spoken utterance.

word “mountain”, it is prompted to identify where in  $\mathbf{x}_{\text{audio}_3}$  it occurs. During VPKL, the model therefore has to first do detection and then localisation, i.e. detection is a task on its own, but localisation includes detection. To do VPKL, we need a multimodal model that can output whether a keyword occurs and at which frame detected keywords occur.

## 3. APPROACH: MULTIMODAL LOCALISATION MODELS

Our VPKL model outputs an overall score  $\mathcal{S} \in [0, 1]$  indicating whether a keyword is present anywhere within an utterance. The keyword is detected if the  $\mathcal{S}$  is above a threshold  $\alpha$ . Additionally, the model outputs a sequence of scores  $\mathbf{a}_{\text{audio}} \in \mathbb{R}^N$  where  $N$  is the number of speech frames; each element  $a_{\text{audio}_i}$  indicates whether the detected keyword occurs at frame  $i$ . To do this, we need a multimodal model that can predict which frames in an utterance is most relevant to a given query image. As starting point for our model, we use the deep audio-visual embedding network (DAVENet) of [4]. We then adapt it by introducing a new sampling scheme and attention mechanism that encourages localisation.

### 3.1. Starting point: DAVENet

DAVENet [4] consists of a vision and an audio network which separately maps an image and its entire spoken caption to single fixed-size embeddings in a common multimodal space. The goal is to get embeddings of paired images and spoken captions to be more similar than the embeddings of mismatched images and captions. Our implementation of DAVENet incorporates some of the extensions from [2].

Following [2], we extend the DAVENet architecture to use ResNet50 [26] for the image network and instead of learning fixed-size embeddings, we learn a sequence of embeddings for each image  $\mathbf{e}_{\text{vision}} \in \mathbb{R}^M$  and caption  $\mathbf{e}_{\text{audio}} \in \mathbb{R}^N$ . Here  $M$  is the number of pixels and  $N$  the number of frames. These embedding sequences are then used in a *matchmap*  $\mathcal{M} \in \mathbb{R}^{M \times N}$  which calculates the dot product between each frame embedding in  $\mathbf{e}_{\text{audio}}$  and each pixel embedding in  $\mathbf{e}_{\text{vision}}$ . The idea is that high similarity in the  $\mathcal{M}$  should indicate those speech frames and image pixels that are related. In [2], the authors showed quantitatively that the matchmaps can indeed localise words and objects corresponding to the same concept.

Another change we make from [4] and [2] is that, instead of using standard speech features as input, we use an acoustic network trained on external data. Concretely, we use a different network as the audio branch in our modified version of DAVENet. This network consists of an acoustic  $f_{\text{acoustic}}$  and a BiLSTM  $f_{\text{BiLSTM}}$  network. For  $f_{\text{acoustic}}$ , we pretrain the contrastive predictive coding (CPC) model of [27] on out-of-domain unlabelled data to obtain more robust acoustic features (by taking advantage of more extensive data sources). The acoustic features are sent to  $f_{\text{BiLSTM}}$  to take advantage of

context from the entire caption. These changes improves the original DAVENet model’s caption retrieval score from 31% to 69% on the English Places corpus [2–4]. We still refer to this custom implementation as DAVENET, since it largely follows the original architecture.

We will see in Section 4 that this multimodal model performs poorly on VPKL. Therefore, we propose two improvements: a keyword sampling scheme (Section 3.2) and a new localising attention mechanism (Section 3.3).

### 3.2. Positive and negative keyword sampling

To learn when an image query occurs in an utterance, we propose a keyword sampling scheme to push image-caption pairs containing the same keyword closer together and pairs not containing the keyword away from each other. We could use an off-the-shelf visual tagger to generate textual tags for the images and use these predicted text labels to sample positive and negative image-caption pairs: if a keyword occurs in the predicted labels of two image-caption pairs, they are positives; if not, they are negatives.

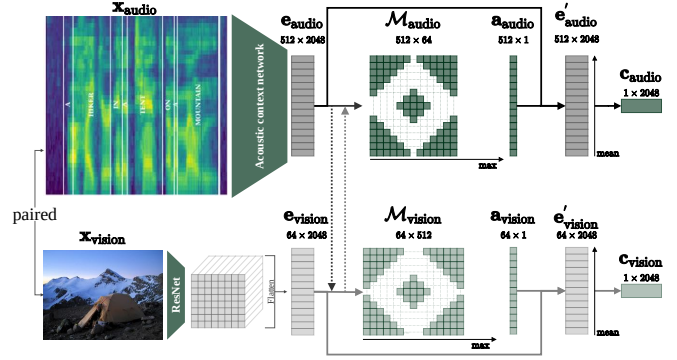
We mainly consider this approach with a perfect tagger, but later in § 4.3 we do initial experiments using a real tagger. In the idealised case we use the ground truth textual transcriptions of the spoken captions to sample positive  $(\mathbf{x}_{\text{audio}}^+, \mathbf{x}_{\text{vision}}^+)$  and three negative  $(\mathbf{x}_{\text{audio}_i}^-, \mathbf{x}_{\text{vision}_i}^-)$  pairs for each image-caption  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  pair. E.g. given an image-caption pair  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  where  $\mathbf{x}_{\text{audio}}$  is “a hiker in a tent on a mountain”, we choose a positive pair also containing the keyword “mountain” like “a mountain with a snowy peak”. Negatives are selected from images that do not contain the keyword “mountain” in the transcription of their captions.

Since we use these pairs in a contrastive loss, we refer to this model as CONTRASTIVEDAVENET. The model is trained to push the similarity scores  $\mathcal{S}$  of pairs containing a certain keyword closer and the similarity scores  $\mathcal{S}$  of pairs not containing the keyword further apart:

$$\begin{aligned} l(\mathbf{e}_{\text{audio}}, \mathbf{e}_{\text{vision}}, \mathbf{e}_{\text{audio}}^+, \mathbf{e}_{\text{vision}}^+, \mathbf{e}_{\text{audio}_{1:3}}^-, \mathbf{e}_{\text{vision}_{1:3}}^-) \\ = 2 \times \text{MSE}(\mathcal{S}(\mathbf{e}_{\text{vision}}, \mathbf{e}_{\text{audio}}), 1) \\ + \text{MSE}(\mathcal{S}(\mathbf{e}_{\text{vision}}, \mathbf{e}_{\text{audio}}^+), 1) + \text{MSE}(\mathcal{S}(\mathbf{e}_{\text{audio}}, \mathbf{e}_{\text{vision}}^+), 1) \\ + \sum_{i=1}^3 \left( \text{MSE}(\mathcal{S}(\mathbf{e}_{\text{vision}}, \mathbf{e}_{\text{audio}_i}^-), -1) + \text{MSE}(\mathcal{S}(\mathbf{e}_{\text{audio}}, \mathbf{e}_{\text{vision}_i}^-), -1) \right), \end{aligned}$$

We obtain  $\mathcal{S}$  by taking the matchmap  $\mathcal{M}$  between the given embedding sequences before max-pooling over the pixel axis and mean-pooling over the temporal axis, as in [2]. We use the mean-squared-error (MSE) between the embeddings with a positive target of 1 and a negative target of  $-1$ . This was based on development experiments: we found that MSE and the wider range of possible target values gave the best results since it pushed the positives and negatives further apart.

Note that CONTRASTIVEDAVENET is the step between DAVENET and our proposed attention model, described next.



**Fig. 3.** The LOCALISATIONATTENTIONNET takes a pair  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  as input. The pair given to the model can also be a positive pair  $(\mathbf{x}_{\text{audio}}^+, \mathbf{x}_{\text{vision}}^+)$  or  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}}^+)$ , or a negative pair  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}_i}^-)$  or  $(\mathbf{x}_{\text{audio}_i}^-, \mathbf{x}_{\text{vision}})$ .

### 3.3. Localising attention mechanism

We make use of an attention mechanism similarly to [9, 10, 12, 28]. These models all take a written keyword as a query, and then look up a word embedding corresponding to that keyword, which is then used to calculate per-frame attention weights  $\mathbf{a}_{\text{audio}}$  over the per-frame audio embeddings  $\mathbf{e}_{\text{audio}}$ . In our case we instead use the matchmap  $\mathcal{M}_{\text{audio}}$ , as shown in Fig. 3. As a reminder, the audio matchmap  $\mathcal{M}_{\text{audio}}$  is calculated by taking the dot product between each audio embedding in  $\mathbf{e}_{\text{audio}}$  and each vision embedding in  $\mathbf{e}_{\text{vision}}$ . We then take the maximum over  $\mathcal{M}_{\text{audio}}$ ’s pixel axis to obtain  $\mathbf{a}_{\text{audio}}$ . Thereafter,  $\mathbf{a}_{\text{audio}}$  is used to weigh  $\mathbf{e}_{\text{audio}}$  to get  $\mathbf{e}'_{\text{audio}}$ . From  $\mathbf{e}'_{\text{audio}}$ , we sum over the temporal axis to get a context vector  $\mathbf{c}_{\text{audio}}$ . In the vision branch we do something very similar: we take the maximum over  $\mathcal{M}_{\text{vision}}$ ’s temporal axis to obtain  $\mathbf{a}_{\text{vision}}$  from which we can calculate  $\mathbf{e}'_{\text{vision}}$  and  $\mathbf{c}_{\text{vision}}$  (similar to how we obtained  $\mathbf{e}'_{\text{audio}}$  and  $\mathbf{c}_{\text{audio}}$ ). It is important to note that  $\mathcal{M}_{\text{vision}}$  is the transpose of  $\mathcal{M}_{\text{audio}}$ .

We train this model, which we refer to as LOCALISATIONATTENTIONNET, by using the  $\mathbf{c}$  vectors in a contrastive loss:

$$\begin{aligned} l(\mathbf{c}_{\text{audio}}, \mathbf{c}_{\text{vision}}, \mathbf{c}_{\text{audio}}^+, \mathbf{c}_{\text{vision}}^+, \mathbf{c}_{\text{audio}_{1:3}}^-, \mathbf{c}_{\text{vision}_{1:3}}^-) \\ = \sum_{i=1}^3 \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{audio}}, \mathbf{c}_{\text{vision}}), 1) + \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{audio}}, \mathbf{c}_{\text{audio}}^+), 1) \\ + \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{audio}}, \mathbf{c}_{\text{audio}_i}^-), -1) + \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{vision}}, \mathbf{c}_{\text{audio}}), 1) \\ + \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{vision}}, \mathbf{c}_{\text{audio}}^+), 1) + \text{MSE}(\mathcal{S}(\mathbf{c}_{\text{vision}}, \mathbf{c}_{\text{vision}_i}^-), -1), \end{aligned}$$

where  $\mathcal{S}$  is the cosine similarity. This loss attempts to push  $\mathbf{c}_{\text{audio}}$  closer to its paired image  $\mathbf{c}_{\text{vision}}$  and to a caption also containing the same keyword  $\mathbf{c}_{\text{audio}}^+$ . At the same time, it should push  $\mathbf{c}_{\text{audio}}$  away from any caption  $\mathbf{c}_{\text{audio}_i}^-$  not containing the keyword. The same goes for the vision part. The idea is to force the model to isolate the keywords common to image-caption pairs and their positive samples.

## 4. EXPERIMENTS

The goal in VPKL is to detect and localise a visual keyword in a spoken utterance. To do this, we train three speech-vision models to predict which (if any) of the frames are related to a given image query.

### 4.1. Experimental setup

**Data:** We pretrain the acoustic network  $f_{\text{acoustic}}$  on the combination of LibriLight [29] and the multilingual Places dataset [5]. We train DAVENET, CONTRASTIVEDAVENET and LOCALISATIONATTENTIONNET on the Flickr8k Audio Captions Corpus [30]. The corpus consists of 8k images where each image is paired with five parallel spoken English captions and is split into train, development and test sets of 30k, 5k and 5k utterances, respectively. The spoken captions are parametrised as mel-spectrograms, with a hop length of 10 ms, a window width of 25 ms and 40 mel bins. We truncate or zero-pad the mel-spectrograms to 1024 frames. Images are resized to  $224 \times 224$  pixels and normalised according to the VGG [31] mean and variances calculated using ImageNet [32].

**Models:** For the image branch, we use the ResNet50 [26] adaptation of [33]. The architecture used for the audio branch of all three models is given in the source code.<sup>1</sup> For training CONTRASTIVEDAVENET and LOCALISATIONATTENTIONNET we use the training set’s transcriptions to sample positive pairs  $(\mathbf{x}_{\text{audio}}^+, \mathbf{x}_{\text{vision}}^+)$  and three negative pairs  $(\mathbf{x}_{\text{audio}_i}^-, \mathbf{x}_{\text{vision}_i}^-)$  for each image-caption  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  pair. For validation on CONTRASTIVEDAVENET and LOCALISATIONATTENTIONNET, we follow the same procedure to sample a positive pair  $(\mathbf{x}_{\text{audio}}^+, \mathbf{x}_{\text{vision}}^+)$  and a negative pair  $(\mathbf{x}_{\text{audio}}^-, \mathbf{x}_{\text{vision}}^-)$  for each image-caption  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  pair using the development set. The validation task measures whether the model will place the positive pair closer to  $(\mathbf{x}_{\text{audio}}, \mathbf{x}_{\text{vision}})$  than it would the negative pair. DAVENET is trained and validated on Flickr8k following a similar setup to the original paper [4]. All models are trained with Adam [34] for 100 epochs using early stopping. The batch sizes and learning rates for each model are tuned on development data.

**Evaluation:** We evaluate our approach on a set of 34 keywords. This is a subset of the keywords from [35], where we only consider keywords that can reasonably be localised using an image. For each keyword, we obtain 10 images from the Flickr8k test split and then manually crop the image region corresponding to the keyword.<sup>2</sup> This cropped image then serves as the image query for that keyword. At

<sup>1</sup>DAVENET: [https://github.com/LeanneNortje/DAVENet\\_VPKL](https://github.com/LeanneNortje/DAVENet_VPKL);

CONTRASTIVEDAVENET: [https://github.com/LeanneNortje/ContrastiveDAVENet\\_VPKL](https://github.com/LeanneNortje/ContrastiveDAVENet_VPKL);

LOCALISATIONATTENTIONNET: [https://github.com/LeanneNortje/LocalisationAttentionNet\\_VKPL](https://github.com/LeanneNortje/LocalisationAttentionNet_VKPL)

<sup>2</sup>VPKL task: <https://github.com/LeanneNortje/VPKL>

test time, we obtain  $\mathbf{c}_{\text{audio}}$  for a test utterance and  $\mathbf{c}_{\text{vision}}$  for an image query. These are compared using cosine similarity. If the score is above a threshold  $\alpha$  for a given caption and image query, the keyword is detected in the caption. The  $\alpha$  for each model is tuned on the development set:  $\alpha_{\text{DAVENET}} = 0.85$ ,  $\alpha_{\text{CONTRASTIVEDAVENET}} = 0.55$  and  $\alpha_{\text{LOCALISATIONATTENTIONNET}} = 0.5885$ . For localisation, if a keyword is detected, then the position of the keyword is taken as the frame where the maximum attention weight occurs. This is compared to ground truth alignments, with a true positive occurring when the predicted frame falls within the time-span of the keyword in the reference. If a keyword is not detected or the prediction does not fall within the time-span of the keyword, this is counted as a mistake. We also implement a random baseline, where a detection score is randomly sampled between  $-1$  and  $1$  and the attention weights for each frame is randomly assigned a value between  $0$  and  $1$ .

**Visual BoW baseline:** Previous work [24, 25] used a visually grounded BoW model to detect written keywords. They use a visual tagger to extract BoW labels for each training image, which then serves as targets for the corresponding spoken caption. The resulting model can be used to detect written keywords in utterances. This was extended to not just detect whether the keyword occurs in the utterance, but also to determine where in the utterance the word occurs (if it was detected) [9, 11, 12]. We compare our models to their model, which takes a written query rather than an image query as input. The visual BoW model is evaluated on the same 34 keywords as our visually prompted models.

### 4.2. Results

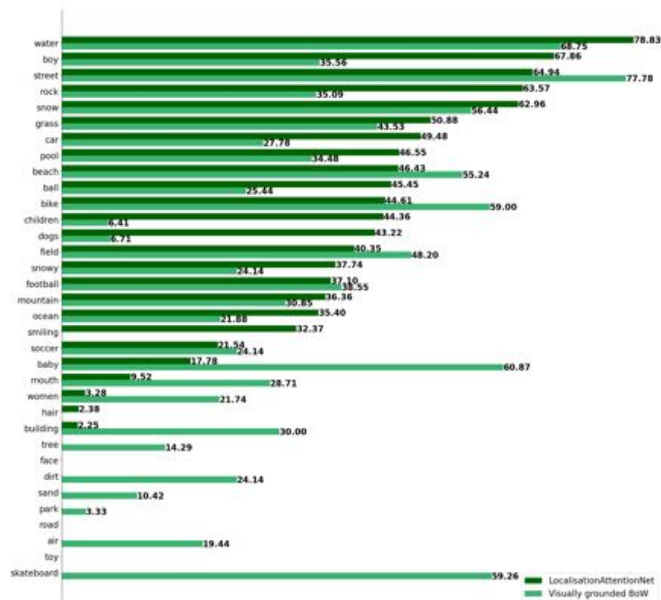
Tables 1 and 2 compare our VPKL models to the visual BoW model for keyword detection and localisation. We see that DAVENET outperforms the random baseline but not the visually grounded BoW model. Moreover, CONTRASTIVEDAVENET outperforms this BoW model on recall and F1, but not in precision. The improvements of CONTRASTIVEDAVENET over DAVENET shows that by sampling positives and negatives using the scheme introduced in Section 3.2, we can better detect and localise visual keywords.

Building on this, LOCALISATIONATTENTIONNET improves

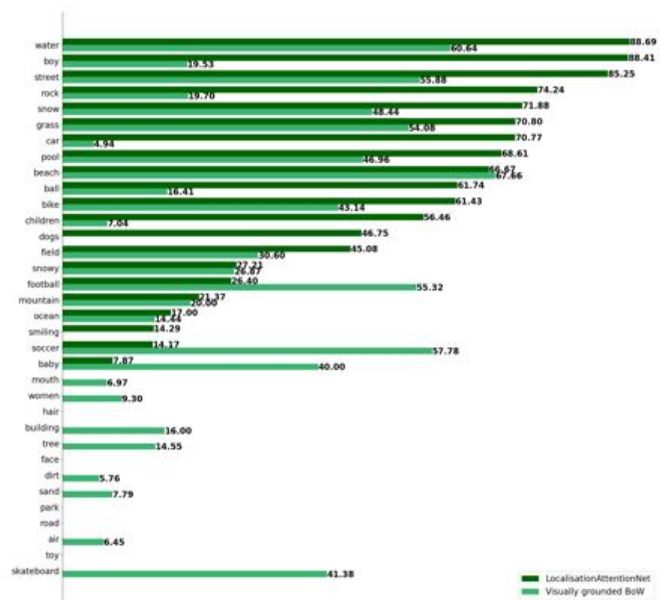
**Table 1.** Keyword detection results (%).

Model	Precision	Recall	F1 score
<i>Text query</i>			
Visually grounded BoW [11]	42.29	36.32	39.08
<i>Image query</i>			
Random baseline	2.30	13.96	3.94
DAVENET	8.86	46.51	14.88
CONTRASTIVEDAVENET	37.97	44.84	41.12
LOCALISATIONATTENTIONNET	<b>48.41</b>	<b>55.85</b>	<b>51.86</b>





a. Detection accuracy scores



b. Localisation F1 scores

**Fig. 4.** The per keyword (a) detection accuracy scores and (b) localisation F1 score achieved by LOCALISATIONATTENTIONNET and the visually grounded BoW model.

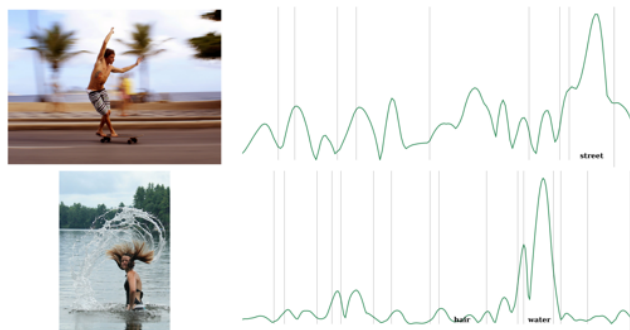
**Table 2.** Keyword localisation results (%).

Model	Precision	Recall	F1 score
<i>Text query</i>			
Visually grounded BoW [11]	33.39	31.02	32.17
<i>Image query</i>			
Random baseline	0.13	0.87	0.22
DAVENET	5.17	33.36	8.95
CONTRASTIVEDAVENET	30.43	39.45	34.36
LOCALISATIONATTENTIONNET	<b>44.43</b>	<b>53.77</b>	<b>48.66</b>

not only the precision and F1 scores of CONTRASTIVEDAVENET, but also the recall. LOCALISATIONATTENTIONNET outperforms the visually grounded BoW model which shows that VPKL can be more accurate than keyword embedding localisation if an appropriate speech-vision model with a localisation objective is used.

### 4.3. Further analyses

The metrics above are based on aggregating scores across keyword types; we now consider per-keyword performance to get a better understanding of how and when the VPKL models outperform the visual BoW model. Fig. 4 shows the (a) detection accuracy scores and (b) localisation F1 scores for each keyword for LOCALISATIONATTENTIONNET and the visual BoW model. We see that LOCALISATIONATTENTION-



**Fig. 5.** The audio attention weights from LOCALISATIONATTENTIONNET for two utterance-query pairs.

NET learns to detect and localise certain keywords more accurately than the BoW model. However, LOCALISATIONATTENTIONNET detects and localises a smaller distribution of keywords, especially when it comes to localisation. E.g., we see worse performance for keywords corresponding to verbs, which makes sense since these are harder to depict with images: “sits”, “sitting”, “riding” and “rides”. Other keywords that the model struggles with are colours (e.g. “orange”) or keywords that are very general (e.g. “air”).

We also see that LOCALISATIONATTENTIONNET can detect some of the keywords but not localise them accurately. In order to see why this happens, we qualitatively look at attention weights produced by the model when given an image query. Fig. 5 shows two examples. In the top example, only one keyword “street” is present, and in the bottom example

**Table 3.** Keyword detection results (%) on the initial 34 keywords.

Model	Precision	Recall	F1 score
Random baseline	2.30	13.96	3.94
DAVENET	8.86	46.51	14.88
LOCALISATIONATTENTIONNET	<b>48.41</b>	<b>55.85</b>	<b>51.86</b>
LOCALISATIONATTENTIONNET trained on generated tags for 34 keywords	31.02	31.83	31.42
LOCALISATIONATTENTIONNET trained on 190 ground truth keywords	20.77	37.06	26.62

**Table 4.** Keyword localisation results (%) on the initial 34 keywords.

Model	Precision	Recall	F1 score
Random baseline	0.13	0.87	0.22
DAVENET	5.17	33.36	8.95
LOCALISATIONATTENTIONNET	<b>44.43</b>	<b>53.77</b>	<b>48.66</b>
LOCALISATIONATTENTIONNET trained on generated tags for 34 keywords	23.20	25.75	24.21
LOCALISATIONATTENTIONNET trained on 190 ground truth keywords	14.48	28.93	19.30

two keywords “water” and “hair” are present in the query. The model can output multiple detection scores for an image query but only one set of attention scores. I.e. we might be able to detect both “water” and “hair”, but we can only localise one keyword. This problem cannot be fixed by taking a smaller crop around *hair* for the keyword “hair”, since the image would still visually contain *water*. Therefore, the model cannot localise two keywords when both are present in an image query.

To investigate the scalability of LOCALISATIONATTENTIONNET, we extend the number of keywords to be learnt to 190.<sup>3</sup> These include the initial 34 keywords. The results in Table 3 and 4 (in line 5) show that extending the number of keywords leads to a decrease in performance on the original 34 keywords. However, our approach still outperforms DAVENET, which has no form of keyword learning.

In moving towards open vocabulary VPCL in zero-resource settings, it is therefore clear that some form of keyword learning would be beneficial. VPCL with an open vocabulary would ideally allow for detection of any keyword depicted using a visual image query. An evaluation on more keywords is one step in this direction, but we are still considering an idealised case using transcriptions to find positive and negative examples. We therefore also present initial experiments where we instead use an actual visual tagger.<sup>4</sup>

<sup>3</sup>LOCALISATIONATTENTIONNET on 190 ground truth keywords: [https://github.com/LeanneNortje/LocalisationAttentionNet\\_VPKL\\_on\\_190\\_keywords](https://github.com/LeanneNortje/LocalisationAttentionNet_VPKL_on_190_keywords)

<sup>4</sup>LOCALISATIONATTENTIONNET on generated tags: [https://github.com/LeanneNortje/LocalisationAttentionNet\\_VPKL\\_on\\_generated\\_tags](https://github.com/LeanneNortje/LocalisationAttentionNet_VPKL_on_generated_tags)

For this approach on LOCALISATIONATTENTIONNET, we use the of-the-shelf visual tagger of [35] to tag the training images with possible keywords. To sample positive and negative image-caption pairs, we then use these predicted keywords. We only use the pairs in which one or more of the 34 keywords (§4.1) occurs to sample positive and negatives: pairs with the same predicted keywords are positives and pairs with different predicted keywords are negatives. From Table 3 and 4, we see that this approach (line 4) outperforms DAVENET (line 2). We conclude that any of our keyword sampling approaches, whether supervised or unsupervised, leads to keyword learning.

Both the generated keyword approach and the 190-word extended keyword evaluations fall short of LOCALISATIONATTENTIONNET trained on ground truth keywords. Therefore, the question remains: how can we extend this approach to work on zero-resource languages? Future work could look into finding a more accurate visual tagger. However, the problem of the model only being able to learn a small amount of keywords remains. We recommend looking into multimodal few-shot learning: learning a new keyword from a few paired images and spoken captions containing the keyword [36]. Recently, [37] looked into few-shot learning using natural images and spoken captions to learn new concepts, but they required a large number of examples to learn a new concept. Although [36, 38, 39] only learned digit classes, they showed that new concepts can be learned for only a few examples. The hope is that few-shot learning would enable LOCALISATIONATTENTIONNET to learn more keywords.

## 5. CONCLUSION

We proposed the new visually prompted keyword localisation (VPKL) task and adapted previous speech-vision models to perform better on this task. Concretely, we proposed a new data sampling scheme, where we use images to find positive utterance pairs containing the same keyword, and we proposed a new localising attention mechanism over matchmaps. We showed that both these innovations gave improvements in VPCL over previous models.

In our main experiments we used the transcriptions of the spoken captions in order to simulate an ideal tagger for our sampling scheme. In further initial experiments, we also adapted the sampling approach to use a real tagger to generate image tags. Although the method using an ideal tagger outperformed the one using a real tagger, both methods outperformed DAVENET, which uses no keywords sampling method. This proves that our sampling approach leads to keyword learning. We also implemented a model with a larger vocabulary size using the ground truth transcriptions. This caused a drop in performance compared to the case where less keywords are learned, but still gave improvements over DAVENET. Future work will look into using few-shot learning to learn more keywords for zero-resource systems.

## 6. REFERENCES

- [1] G. Chrupała, “Visually grounded models of spoken language: A survey of datasets, architectures and evaluation techniques,” *J. Artif. Intell. Res.*, vol. 73, 2022.
- [2] 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.
- [3] D. Harwath and J. R. Glass, “Learning word-like units from joint audio-visual analysis,” *arXiv preprint arXiv:1701.07481*, 2017.
- [4] D. Harwath, A. Torralba, and J. Glass, “Unsupervised learning of spoken language with visual context,” in *Proc. NIPS*, 2016.
- [5] D. Harwath, G. Chuang, and J. Glass, “Vision as an Interlingua: Learning multilingual semantic embeddings of untranscribed speech,” in *Proc. ICASSP*, 2018.
- [6] P. Peng and D. Harwath, “Word discovery in visually grounded, self-supervised speech models,” in *Proc. Interspeech*, 2022.
- [7] —, “Fast-slow transformer for visually grounding speech,” in *Proc. ICASSP*, 2022.
- [8] 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.
- [9] K. Olaleye and H. Kamper, “Attention-based keyword localisation in speech using visual grounding,” in *Proc. Interspeech*, 2021.
- [10] N. C. Tamer and M. Saraçlar, “Keyword search for sign language,” in *Proc. ICASSP*, 2020.
- [11] K. Olaleye, B. van Niekerk, and H. Kamper, “Towards localisation of keywords in speech using weak supervision,” in *Proc. NeurIPS-SAS*, 2020.
- [12] K. Olaleye, D. Oneata, and H. Kamper, “Keyword localisation in untranscribed speech using visually grounded speech models,” *IEEE J. Sel. Top. Signal Process.*, vol. 16, 2022.
- [13] P. C. Bomba and E. R. Siqueland, “The nature and structure of infant form categories,” *J. Exp. Child Psychol.*, 1983.
- [14] S. Pinker, *The Language Instinct*. New York: Harper Perennial, 1994.
- [15] P. D. Eimas and P. C. Quinn, “Studies on the formation of perceptually based basic-level categories in young infants,” *Child Develop.*, 1994.
- [16] D. Roy, “Grounded spoken language acquisition: experiments in word learning,” *IEEE Trans. Multimedia*, 2003.
- [17] L. Boves, L. ten Bosch, and R. Moore, “ACORNS - towards computational modeling of communication and recognition skills,” in *Proc. ICCI*, 2007.
- [18] L. Gelderloos and G. Chrupała, “From phonemes to images: levels of representation in a recurrent neural model of visually-grounded language learning,” in *Proc. COLING*, 2016.
- [19] O. Räsänen and K. Khorrami, “A computational model of early language acquisition from audiovisual experiences of young infants,” in *Proc. Interspeech*, 2019.
- [20] G. A. Miller and P. M. Gildea, “How children learn words,” *SciAM*, vol. 257, 1987.
- [21] A. Garcia and H. Gish, “Keyword spotting of arbitrary words using minimal speech resources,” in *Proc. ICASSP*, 2006.
- [22] I. Szoke, P. Schwarz, P. Matejka, L. Burget, M. Karafiat, M. Fapso, and J. Cernocky, “Comparison of keyword spotting approaches for informal continuous speech,” in *Proc. Interspeech*, 2005.
- [23] J. Wilpon, L. Rabiner, C.-H. Lee, and E. Goldman, “Automatic recognition of keywords in unconstrained speech using hidden Markov models,” *IEEE Trans. Acoust. Speech Signal Process.*, vol. 38, 1990.
- [24] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” in *Proc. Interspeech*, 2017.
- [25] H. Kamper and M. Roth, “Visually grounded cross-lingual keyword spotting in speech,” in *Proc. SLTU*, 2018.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016.
- [27] 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.
- [28] D. Palaz, G. Synnaeve, and R. Collobert, “Jointly learning to locate and classify words using convolutional networks,” in *Proc. Interspeech*, 2016.

- [29] J. Kahn, M. Riviere, W. Zheng, E. Kharitonov, Q. Xu, P. Mazare, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, “Libri-light: A Benchmark for ASR with limited or no supervision,” in *Proc. ICASSP*, 2020.
- [30] D. Harwath and J. Glass, “Deep multimodal semantic embeddings for speech and images,” in *Proc. ASRU*, 2015.
- [31] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *Proc. ICLR*, 2015.
- [32] 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.
- [33] D. Harwath, W.-N. Hsu, and J. Glass, “Learning hierarchical discrete linguistic units from visually-grounded speech,” *arXv preprint arXiv:1911.09602*, 2020.
- [34] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.
- [35] 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.*, vol. 27, 2019.
- [36] R. Eloff, H. A. Engelbrecht, and H. Kamper, “Multimodal one-shot learning of speech and images,” in *Proc. ICASSP*, 2019.
- [37] T. Miller and D. Harwath, Dawid, “Exploring few-shot fine-tuning strategies for models of visually grounded speech,” in *Proc. Interspeech*, 2022.
- [38] L. Nortje and H. Kamper, “Unsupervised vs. transfer learning for multimodal one-shot matching of speech and images,” in *Proc. Interspeech*, 2020.
- [39] —, “Direct multimodal few-shot learning of speech and images,” in *Proc. Interspeech*, 2021.