

Visual Keyword Spotting with Attention

K R Prajwal*
prajwal@robots.ox.ac.uk

Liliane Momeni*
liliane@robots.ox.ac.uk

Triantafyllos Afouras
afourast@robots.ox.ac.uk

Andrew Zisserman
az@robots.ox.ac.uk

Visual Geometry Group
Department of Engineering Science
University of Oxford
Oxford, UK

Abstract

In this paper, we consider the task of spotting spoken keywords in silent video sequences – also known as visual keyword spotting. To this end, we investigate Transformer-based models that ingest two streams, a visual encoding of the video and a phonetic encoding of the keyword, and output the temporal location of the keyword if present. Our contributions are as follows: (1) We propose a novel architecture, the *Transpotter*, that uses full cross-modal attention between the visual and phonetic streams; (2) We show through extensive evaluations that our model outperforms the prior state-of-the-art visual keyword spotting and lip reading methods on the challenging LRW, LRS2, LRS3 datasets by a large margin; (3) We demonstrate the ability of our model to spot words under the extreme conditions of isolated mouthings in sign language videos.

1 Introduction

In recent years, there has been significant progress in automatic visual speech recognition (VSR) due to the availability of large-scale annotated datasets and the development of powerful neural network-based learners [0, 0, 13]. These methods are continually improving and becoming more sophisticated, by incorporating better visual models, stronger language modelling and training on larger datasets. Indeed the best industrial grade lip reading models today are far superior to humans, and achieve error rates approaching Automatic Speech Recognition (ASR) performance [32, 41].

However, for many applications it is not necessary to transcribe every word that is spoken in a silent video (the task of VSR), rather only specific utterances or keywords need to be recognised. This is for example the case in “wake word” recognition, where only particular keywords need to be spotted over long input sequences. A further drawback of VSR methods is that they are heavily reliant on language modelling; in general, their performance decreases significantly when context is limited (e.g. short utterances) or parts of the input are occluded, e.g. from the speaker’s hands or a microphone. In this work, we focus instead on the task of

Visual Keyword Spotting (KWS), where the goal is to detect and localise a *given* keyword in (silent) spoken videos.

Automatic visual KWS enables a diverse range of practical applications: indexing archival silent videos by keyword to enable content-based search; helping virtual assistants (e.g. Alexa and Siri) and smart home technologies respond to wake words and phrases; assisting people with speech impairment (e.g. amyotrophic lateral sclerosis patients) or aphonia in communication [56]; and detecting mouthings in sign language videos [9].

KWS differs in complexity from VSR primarily because in KWS we are armed with the keyword we need to recognise, whereas VSR has the harder task of recognising every word from scratch. The core hypothesis motivating this work is that *this additional knowledge renders visual KWS an easier task* than VSR; and it is therefore expected that KWS should achieve a higher performance than VSR, and generally be more robust to challenging and adversarial situations. Nevertheless, visual KWS remains a *very difficult* task and shares similar challenges to VSR methods: first, some words sound different but involve identical lip movements ('man', 'pan', 'ban'), these *homopheme* words cannot be distinguished using only visual information. Second, speech variations such as accents, speed, and mumbling can alter lip movements significantly for the same word. Third, co-articulation of the lips between preceding and subsequent words in continuous speech also affects lip appearance and motion.

In this paper, we make the following three contributions: (i) We propose a novel Transformer-based architecture, the *Transpotter* (a portmanteau of *Transformer* and *Spotter*), that is tailored to the visual KWS task. The model takes as input two streams, one encoding visual information from a video and the other providing a phonetic encoding of the keyword; the heterogeneous inputs are then fused using full cross-modal attention. (ii) Through extensive evaluations, we show that our Transpotter model outperforms the prior state-of-the-art visual KWS and VSR methods on the challenging LRW, LRS2 and LRS3 lip reading datasets by a large margin. (iii) We test our best model under extreme conditions: finding words in mouthings of people communicating using sign language. Signers sometimes mouth words as they sign as an additional non-manual signal to disambiguate and help understanding [60]. This new task is extremely challenging as there is a significant domain shift between full spoken sentences (in our training and test sets) and mouthings, where the context is sporadic and phonemes of the keyword may be missing – as sometimes only parts of words are mouthed [10]. Our approach outperforms previous KWS models in this challenging, practical use-case. Video examples are available at the project's webpage: www.robots.ox.ac.uk/~vgg/research/transpotter.

2 Related work

Our work relates to prior work on KWS, lip reading, visual grounding, and applications of Transformers for text and video. We present a brief discussion of these topics below.

KWS. KWS in audio (speech) is a well studied problem with a long history, spanning several decades. Prior to the establishment of deep learning models, KWS methods were based on Hidden Markov Models [50, 65], dynamic time warping [62, 63, 75] or indexing of ASR lattices [13]. A number of works have since used deep architectures suitable for sequence modelling (e.g. RNNs, CNNs, or graph convolutional networks) [6, 15, 23, 31, 36, 38, 45, 52, 60, 64, 76], including encoder-decoder approaches [8, 51, 73, 78]. Berg *et al.* [9] recently proposed using a Transformer model for the same task. Different from ours, this work uses a single input stream (audio) and only learns to spot a fixed vocabulary of keywords. In contrast, we use Transformers to temporally process, then fuse the multi-modal inputs, building a

model that can eventually perform open-set KWS. Visual KWS has also received attention recently. The proposed methods include query-by-example [63] approaches, sliding window classification [69], or looking up phonetic queries in lip reading feature sequences [43, 69], while audio-visual methods [20, 43, 62] that fuse the two modalities to improve robustness to noise have also been proposed. Our method builds upon these approaches: we address various weaknesses and propose superior video-text modelling as well as explicit keyword localization, resulting in significantly improved performance.

Lip reading. Early works in lip reading usually relied on hand-crafted pipelines and features [26, 46, 48, 77]. The availability of large scale lip reading datasets [0, 18] and the development of deep neural network models resulted in major performance improvements, initially in word-level lip reading [16, 58] and constrained sentences [0]. Sentence level models were subsequently developed, using sequence-to-sequence architectures based on RNNs [18], CTC-based [56] approaches, or a hybrid of the two [47]. Replacing RNNs with Transformers resulted in better performing architectures [0, 27, 74]. Joint audio-visual training and cross-modal distillation [9, 39, 70] have also been investigated. The current state-of-the-art model uses Transformers in the visual front-end and achieves remarkable results with word error rates reaching as low as 30.7% [54].

Visual grounding. Our work is also related to tasks such as natural language grounding in videos [14, 24, 25, 29, 30, 68, 71, 72] and subtitle alignment in sign language clips [17].

Transformers. Since their introduction for machine translation, Transformers [63] have become ubiquitous and are used today in a wide range of applications from natural language processing [19, 49] and speech recognition [21, 65, 42] to visual representation learning [10, 22, 34, 66]. In this work, we rely on Transformers as our building blocks for their strong sequence modelling capability and inherent potential for localisation through attention.

3 Visual KWS with Attention

In this section, we describe our proposed method shown in Figure 1. We outline the architecture of our model (Section 3.1), our training procedure (Section 3.2) and differences to prior work (Section 3.3). We refer the reader to the Appendix for further details.

3.1 The Transpotter Architecture

Our model ingests two input streams: (i) a textual keyword $q = (q_1, q_2 \dots, q_{n_p})$, and (ii) a silent video clip $v \in \mathbb{R}^{T \times H \times W \times 3}$ in which we need to spot the keyword. For each of the inputs, we have separate encoders that learn initial modality-specific representations. This is followed by a joint multi-modal Transformer that learns cross-modal relationships between the video and text features. The joint transformer predicts two outputs: (i) a sequence-level probability of the keyword occurring in the video and (ii) frame-level probabilities indicating the location of the keyword in the video if present. We describe each of the modules next.

Text Representations. Our textual input is a phonetic representation of the keyword, obtained using a pronunciation dictionary. The input phoneme sequence of length n_p is mapped to a sequence of learnable embedding vectors $Q \in \mathbb{R}^{n_p \times d}$. Sinusoidal positional encodings are added to the input phoneme feature vectors, and the result is passed through a Transformer Encoder [63] with N_t layers to capture temporal information across the phoneme sequence:

$$Q_{enc} = \text{encoder}_q(Q + PE_{1:n_p}) \in \mathbb{R}^{n_p \times d}.$$

Video Representations. We use a pre-trained visual front-end (either a CNN [0] or VTP [54]) to extract a feature vector for each input video frame, $V \in \mathbb{R}^{T \times d}$. Similar to the text encoder

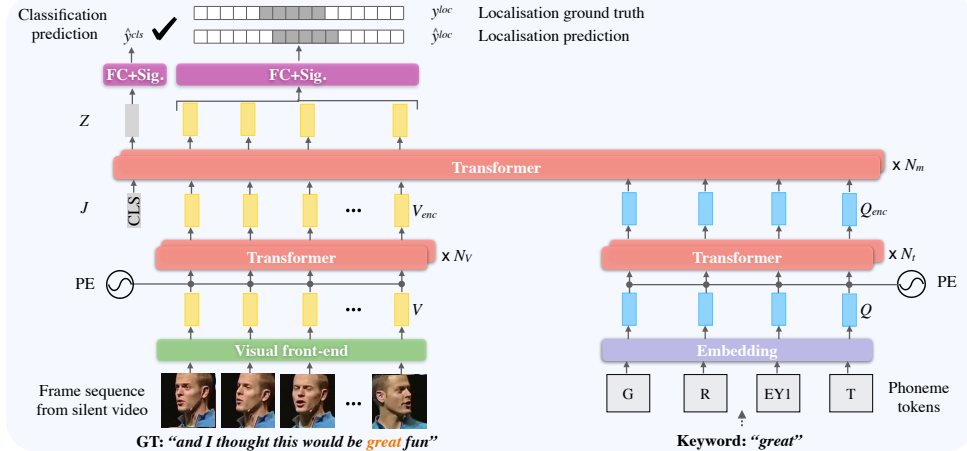


Figure 1: **The Transpotter architecture:** Video frames are inputted to a visual front-end (CNN [10] or VTP [32]) to extract low-level visual features, which are then passed to N_v Transformer layers to encode temporal information. The keyword in the form of a phoneme sequence is encoded using N_t Transformer layers. The text and visual features are finally concatenated in time and processed using a joint multi-modal Transformer which predicts: (i) the probability the keyword occurs in the video, (ii) frame-level probabilities indicating the location of the word. PE corresponds to positional encoding.

Q_{enc} , we pass the visual features through a Transformer Encoder with N_v layers to capture temporal information, after adding positional encodings:

$$V_{enc} = \text{encoder}_v(V + PE_{1:T}) \in \mathbb{R}^{T \times d}.$$

Joint Video-Text Representations. The uni-modal representations V and Q are concatenated along the time dimension to produce a single sequence of feature vectors. A learnable $[CLS]$ token embedding (such as in BERT [19] and ViT [2]) is then prepended to the result:

$$J = ([CLS]; V_{enc}; Q_{enc}) \in \mathbb{R}^{(1+T+n_p) \times d}.$$

We use a Transformer encoder with N_m layers to jointly learn the relationships across video and phoneme vectors:

$$Z = \text{encoder}_{vq}(J + PE_{1:(1+T+n_p)}) \in \mathbb{R}^{(1+T) \times d}.^1$$

Prediction heads. The $[CLS]$ output feature vector Z_1 serves as a joint aggregate representation for the video-text pair. An MLP head for binary classification, f_c is attached to Z_1 to predict the probability of the keyword being present in the video:

$$\hat{y}^{cls} = \sigma(f_c(Z_1)) \in \mathbb{R}^1,$$

where σ denotes a sigmoid activation. To localise the keyword, we attach a second MLP head f_l that is shared across all the video output states from the multi-modal joint Transformer:

$$\hat{y}^{loc} = \sigma(f_l(Z_{2:(T+1)})) \in \mathbb{R}^T.$$

The output y_t^{loc} at each video frame time-step $t \in T$ indicates the probability of the frame t being a part of the keyword utterance.

¹ the n_p outputs corresponding to the phonetic embeddings are dropped.

3.2 Training

Optimisation objectives. Given a training dataset \mathcal{D} consisting of tuples (v, q, y^{cls}, y^{loc}) of silent video clips, text queries, class labels and location labels (indicating the position of the keyword within the clip), we define the following objectives:

$$\mathcal{L}^{cls} = -\mathbb{E}_{(v, q, y^{cls}) \in \mathcal{D}} BCE(y^{cls}, \hat{y}^{cls}) \quad (1)$$

$$\mathcal{L}^{loc} = -\mathbb{E}_{(v, q, y^{cls}, y^{loc}) \in \mathcal{D}} y^{cls} \left[\frac{1}{T} \sum_{t=1}^T BCE(y_t^{loc}, \hat{y}_t^{loc}) \right] \quad (2)$$

$$BCE(y, \hat{y}) = y \log \hat{y} + (1 - y) \log(1 - \hat{y}), \quad (3)$$

where BCE stands for the binary cross-entropy loss. The labels y^{cls} are set to 1 when the given keyword occurs in the video and 0 otherwise; the frame labels y^{loc} are set to 1 for the frames where the keyword is uttered and 0 otherwise. We train the model to optimise the total loss $\mathcal{L} = \lambda \mathcal{L}^{cls} + (1 - \lambda) \mathcal{L}^{loc}$, where λ is a balancing hyper-parameter.

3.3 Discussion

Compared to prior approaches, the design of our model offers several important advantages.

Stronger Visual Representations. Previous works [43, 59] model temporal relationships between video frames using RNNs. In contrast, we employ Transformers [63], which are far more effective in modeling temporal relationships [9, 30].

Joint Video-text Modeling. Prior works such as KWS-Net [43] follow a late-fusion strategy. In our model each frame-wise video feature can attend to any keyword token (phoneme) and vice-versa. The information exchange across the modalities occurs at every layer, without restrictions on the receptive field for either modality.

Stronger keyword localisation. Fine-grained localisation of the keyword in the video can be important for applications such as sign spotting [6]. Existing methods [43, 59] “weakly” localise the keyword by taking the sequence-level prediction to be the maximum probability over all the video time-steps. We instead provide stronger frame-level supervision, by enforcing the model to predict the exact temporal extent of the keyword in the video.

3.4 Implementation details

Pre-training the visual front-end. We explore two different visual front-end architectures for the Transpotter: (1) a CNN, highly similar in architecture to TM-seq2seq [9] and (2) VTP [54], the current state-of-the-art for lip reading (trained only on public data). Both models are trained end-to-end on two-word video clips of LRS2 [18] and LRS3 [11] for lip reading. We refer the reader to the Appendix for the exact CNN architecture and training hyper-parameters. We refer the reader to [54] for architectural hyper-parameters and training protocols for VTP. We pre-compute the visual features for each backbone for both datasets and then train directly on them for faster training. All our models and ablations use the pre-trained CNN features, unless otherwise stated.

Sampling. We form the training dataset \mathcal{D} by randomly sampling with 50% probability a positive or negative video clip v for each query q . Each video v contains word boundary annotations, which allows (i) performing data augmentation by randomly cropping video segments during training, and (ii) creating frame labels y^{loc} , as described in 3.2.

Misc. The keyword q is mapped to a phoneme sequence using the CMU dictionary [62]; words not present in the dictionary are discarded from training \mathcal{D} . We set $\lambda = 0.5$.

4 Experiments

This section is structured as follows: We first present the datasets used as well the evaluation protocols that we follow in our experiments (Section 4.1). Next, we compare the performance of our proposed Transpotter model against strong baselines (Section 4.2) and then present a comprehensive study ablating our design choices (Section 4.3). Finally, we perform further performance analysis and provide qualitative results (Section 4.4).

4.1 Datasets and Evaluation Protocol

Datasets. All models and baselines are trained and evaluated on LRS2 [18] and LRS3 [10] lip reading datasets. LRS2 contains BBC broadcast footage from British television and LRS3 is based on TED/TEDx videos downloaded from YouTube (refer to the Appendix for detailed statistics). The video clips for both datasets are tightly cropped face-tracks of active speakers only. For each clip, a full transcription of the utterance as well as word boundary alignments are provided. The number of videos, number of keyword instances and keyword vocabulary for each of the test sets is shown in Table 1.

Evaluation Protocol. Evaluation is performed for every test dataset as follows: First, the vocabulary of test keywords is determined, by considering all the words occurring in the test set transcriptions with above a certain phoneme length n_p . If not specified, we use $n_p \geq 3$. Every word in the query vocabulary is then searched for in all the test set videos.

Metrics. Given ground truth video-keyword samples, we assess the performance of our model in two ways. First, we assess classification performance, *i.e.* whether the model can accurately predict whether the keyword occurs in the video or not. We compute accuracy ($\text{Acc}_{@k}^{Cls}$) and mean average precision (mAP^{Cls}) metrics, where $\text{Acc}_{@k}^{Cls}$ measures how often a given keyword occurs in any of the top- k retrievals, and mAP^{Cls} is obtained with the above criterion (where every word in the test keyword vocabulary is considered as a separate class).

Second, we assess the model’s localisation capability, *i.e.* whether the model can accurately localise the keyword in the video clip. We follow common practice from the detection literature: we consider a keyword accurately detected when the intersection-over-union (IOU) between the prediction \hat{y}^{loc} and ground truth label y^{loc} is above a certain threshold, and calculate the mean average precision mAP^{Loc} . To calculate the IOU, we binarise the model’s predictions using a threshold $\tau = 0.5$.

4.2 Comparison to baselines

We compare our model’s performance against a state-of-the-art VSR model and KWS-Net [13], the previous state-of-the-art visual KWS model.

VSR baseline. We use an improved version of the TM-seq2seq [10] VSR model, with the same pre-trained CNN backbone (Section 3.4) that we use for the KWS models. The model is trained with the curriculum training strategy of [10] (details in the Appendix). The VSR model achieves state-of-the-art Word Error Rate (WER) performance of 36.9% and 48.0% on the LRS2, LRS3 test sets respectively. Since the VSR model only produces text transcriptions of a given video, but no localisation prediction, we can only evaluate its classification performance ($\text{Acc}_{@k}^{Cls}, \text{mAP}^{Cls}$). We follow the method detailed in [18] to estimate the posterior probability that the keyword occurs in a video clip.

KWS-Net. As a KWS baseline we use the state-of-the-art model of Momeni *et al.* [13]. For fair comparison, here too we use the same CNN backbone that is also used for our model.

State-of-the-art KWS. We report our model’s performance and compare it with strong baselines in Table 1. It is clear that our model outperforms both baselines. On the last row,

Model	LRS2				LRS3			
	1.2K vids. / 4.3K inst. / 1.6K vocab.				1.3K vids. / 6.1K inst. / 1.9K vocab.			
	$\text{Acc}_{@1}^{Cls}$	$\text{Acc}_{@5}^{Cls}$	mAP^{Cls}	mAP^{Loc}	$\text{Acc}_{@1}^{Cls}$	$\text{Acc}_{@5}^{Cls}$	mAP^{Cls}	mAP^{Loc}
KWS-Net [43]	36.1	61.2	41.0	36.2	29.8	54.6	34.3	29.2
VSR	63.7	76.3	64.3	-	52.3	66.0	50.3	-
Transpotter	65.0	87.1	69.2	68.3	52.0	77.1	55.4	53.6
Transpotter (VTP)	68.7	90.7	72.5	71.6	55.7	78.5	58.2	56.1

Table 1: Comparison to baselines: We outperform the current state-of-the-art KWS and VSR methods by a large margin. Our Transpotter model is particularly effective in localising the keyword in the video. Moreover, by using the recently proposed VTP [62] architecture as the Transpotter’s visual backbone instead of a CNN, we achieve even better performance.

we show the boost in performance by replacing the CNN with the recently proposed VTP backbone [62], resulting in state-of-the-art performance on both the LRS2 and LRS3 datasets.

Evaluation on LRW. We also compare the performance of KWS-Net [43] with our proposed Transpotter model on the LRW [16] test set following the same evaluation protocol. The test set contains 25K single-word video clips spanning a vocabulary of 500 words (50 instances per word). Note that KWS-Net has been pretrained on the LRW training split, but the Transpotter has only been trained on LRS2 and LRS3. As we can see in Table 2, the Transpotter outperforms the previous state-of-the-art baseline KWS-Net by a large margin. We refer the reader to the Appendix for a qualitative error analysis in this setting.

Model	$\text{Acc}_{@1}^{Cls}$	$\text{Acc}_{@5}^{Cls}$	mAP^{Cls}
KWS-Net [43]	66.6	89.0	33.0
Transpotter	85.8	99.6	64.1

Table 2: Comparison on LRW [16]: The Transpotter outperforms the previous state-of-the-art KWS model on the LRW test set, despite not having been trained on LRW data. The localization metric mAP^{Loc} is not reported as the input videos are single-word clips.

4.3 Architecture ablations

To assess our design choices for the Transformer skeleton, we perform a number of ablations considering variations of the model architecture. We briefly explain the alternative approaches below; more details can be found in the Appendix.

In particular we consider two alternative encoder-decoder architectures, with the video input at the encoder side and the text query at the decoder ($\text{Enc}_{vid}\text{-Dec}_{text}$) and vice versa ($\text{Enc}_{text}\text{-Dec}_{vid}$). Since the latter model outputs at the temporal resolution of the video input, it can explicitly localise the keyword (in the same way as the Transpotter), while the former can only perform classification. We also consider a variant of the Transpotter, where the model does not output localisation predictions (hence no \mathcal{L}^{loc} is used for its training). We show the results in Table 3. The selected Transpotter architecture outperforms all variants. In particular, by comparing rows 2 and 4, we observe that training with a localisation head and loss \mathcal{L}^{loc} also improves classification (e.g. 64.0 vs 69.2 mAP^{Cls}).

4.4 Transpotter performance analysis

In this section, we analyse the performance of our proposed method when varying the keyword length and the size of the surrounding visual context.

Keyword length. In Figure 2a, we plot the model’s performance on the LRS2 test set against the minimum keyword length in phonemes n_p . As expected, longer keywords are easier

Model	LRS2				LRS3			
	Acc $_{@1}^{Cls}$	Acc $_{@5}^{Cls}$	mAP Cls	mAP Loc	Acc $_{@1}^{Cls}$	Acc $_{@5}^{Cls}$	mAP Cls	mAP Loc
Enc $_{vid}$ -Dec $_{text}$	52.5	80.0	57.9	-	40.3	66.9	43.2	-
Transpotter w/o loc.	59.4	84.1	64.0	-	46.5	72.1	49.8	-
Enc $_{text}$ -Dec $_{vid}$	63.8	86.8	68.4	67.8	52.1	76.6	54.9	53.1
Transpotter	65.0	87.1	69.2	68.3	52.0	77.1	55.4	53.6

Table 3: **Model ablations:** Our approach of jointly modeling text and video sequences with a localisation head for stronger supervision outperforms other architectural designs.




Figure 2: (a) Transpotter’s performance increases with the keyword length; (b) Transpotter performs far better than VSR with limited context. Both methods improve with more context.

to spot and therefore result in better retrieval performance. Indeed for long 7-phoneme keywords, mAP Loc reaches as high as 82.5. We note however that even for very challenging short keywords with only 2 phonemes (such as "my", "to", "at"), mAP Loc stays high at 67.5.

Context. The visual appearance of spoken words can be highly ambiguous [2], therefore recognising isolated words from visual input alone may be very challenging. Current lip reading models utilise the surrounding visual context to resolve this ambiguity. In Figure 2b, we illustrate how the performances of our Transpotter KWS model and our VSR baseline vary based on the amount of contextual information available. We plot the mAP Cls against the number of words in the video clip. We observe that both models benefit from larger surrounding context, with the Transpotter outperforming the VSR baseline consistently.

Qualitative analysis. In Figure 3, we show qualitative examples from the LRS2 and LRS3 test sets. It is clear that the model produces smooth predictions that precisely indicate the full location of the word. In the bottom right corner we observe a failure case where the model’s confidence is low – the keyword “that’s” in this case is short.

Model response to homophemes. We further probe our Transpotter model for failure cases. In visual-only keyword spotting, a common failure case is due to homophemes, *i.e.* words with identical lip movements. To investigate the response of our model to such cases, we construct a list of keywords from the LRS2 test set sentences that are known to have homopheme counterparts (*e.g.* *mark*, which has two matching homophemes, *bark* and *park*) and then for each test set clip that contains one of the keywords, we query that keyword along with its corresponding homophemes and plot the model’s outputs. We illustrate several examples in Figure 4. We observe that in such cases, the model spots the keyword as well as its homophemes at the same (ground truth) location.

5 Mouthing Spotting in Sign Language videos

In this section, we investigate the application of our method for spotting mouthed words in sign language videos. This is an important application of visual KWS, as it has enabled an entire line of work on sign language recognition [8, 24, 62].



Figure 3: **Qualitative results on LRS2 and LRS3:** The Transpotter accurately localises the keyword in most examples. In the bottom right example, the model’s confidence is low, most likely because it is a short word. The IOU is zero since we threshold at $\tau = 0.5$.

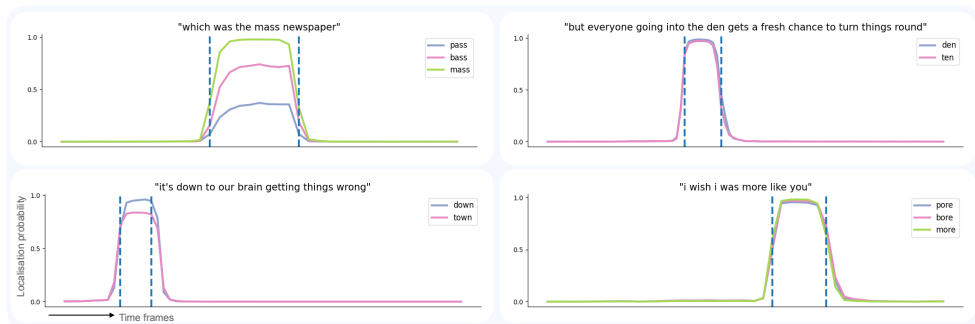


Figure 4: **Model’s response to homophemes:** We query words and their corresponding homophemes for LRS2 test set clips. We observe that the model spots the words and their homophemes at the same (ground truth) location.

Data description & evaluation protocol. Here, we use a subset of BSL Corpus [64, 65] as a test set. BSL Corpus is a large public dataset containing videos of conversations conducted in sign language by deaf signers, from various regions across the UK. We extend the dataset’s annotations by adding a *Mouthing* tier and asking a deaf annotator to identify and localise mouthing occurrences that correspond to visible signs. We obtain 383 mouthing instances, from 29 different signers, over a keyword vocabulary size of 187. We use a pre-processing pipeline similar to [18] to obtain face-cropped tracks around the faces of the signers. To evaluate KWS performance, we take 8-second video clips centered around the annotated mouthings and follow the same evaluation protocol described in Section 4.

Results. We summarise the evaluation results in Table 4. The Transpotter model is far superior to the prior state-of-the-art KWS baseline, achieving a great improvement in performance (e.g. 29.6 vs 15.6 mAP^{Cls} score). To complete this analysis, we also show qualitative examples of the spotted mouthings in Figure 5.

Discussion. We note that sign language mouthings are often very different from equivalent

spoken words. Words may be partially mouthed and can be occluded by the signing hands. There is therefore a significant domain gap between the BSL-Corpus signing videos and our lip reading training videos. However, we note that our proposed model greatly outperforms the KWS-Net baseline – a variant of which has been successfully deployed for detecting mouthings in order to bootstrap learning of sign spotting methods [6, 44, 67]. This indicates the potential of our proposed method to greatly improve these pipelines.

Model	Acc Cls _{@1}	Acc Cls _{@5}	mAP Cls
KWS-Net [43]	12.4	29.6	15.6
Transpotter	22.5	47.1	29.6

Table 4: **Spotting mouthings in BSL-Corpus:** The Transpotter is far more accurate than the current state-of-the-art in spotting keywords in videos.




Figure 5: **Qualitative results on BSL-Corpus:** Despite the large domain shift from our training examples and additional challenges such as partially mouthed words and hand occlusions, the Transpotter succeeds in correctly spotting mouthings in these challenging conditions. We observe a failure case (bottom right) where the localisation is incorrect. We note that contrary to LRS2 and LRS3, where word boundaries are obtained through robust audio-based forced alignment, the annotations for BSL-Corpus are noisier as they are performed manually.

6 Conclusion

We have presented the *Transpotter*, a cross-modal attention based architecture for visual keyword spotting. Our method surpasses the performance of the previous best visual keyword spotting approach by a large margin, as well as that of a state-of-the-art lip reading baseline. We demonstrate the ability of our model to generalise to sign language videos where it can be used to spot mouthings, enabling automatic annotation of sign instances. In future work, we plan to further improve our method’s performance by incorporating keyword semantics and context of the surrounding words.

Acknowledgements. Funding for this research is provided by the UK EPSRC CDT in Autonomous Intelligent Machines and Systems, the Oxford-Google DeepMind Graduate Scholarship, the EPSRC Programme Grant VisualAI (EP/T028572/1) and the Royal Society Research Professorships 2019 RP/R1/191132. We thank Samuel Albanie for his invaluable help in applying our method to signer mouthings.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. LRS3-TED: a large-scale dataset for visual speech recognition. In *arXiv preprint arXiv:1809.00496*, 2018.
- [2] Triantafyllos Afouras, Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Deep audio-visual speech recognition. *IEEE PAMI*, 2019.
- [3] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Asr is all you need: Cross-modal distillation for lip reading. In *International Conference on Acoustics, Speech, and Signal Processing*, 2020.
- [4] Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. Character-level language modeling with deeper self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3159–3166, 2019.
- [5] Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Joon Son Chung, Neil Fox, and Andrew Zisserman. BSL-1K: Scaling up co-articulated sign language recognition using mouthing cues. In *ECCV*, 2020.
- [6] Sercan Arik, Markus Kliegl, Rewon Child, Joel Hestness, Andrew Gibiansky, Chris Fougner, Ryan Prenger, and Adam Coates. Convolutional recurrent neural networks for small-footprint keyword spotting. In *INTERSPEECH*, 2017.
- [7] Yannis M. Assael, Brendan Shillingford, Shimon Whiteson, and Nando de Freitas. Lipnet: Sentence-level lipreading. *arXiv:1611.01599*, 2016.
- [8] Kartik Audhkhasi, Andrew Rosenberg, Abhinav Sethy, Bhuvana Ramabhadran, and Brian Kingsbury. End-to-end asr-free keyword search from speech. *IEEE Journal of Selected Topics in Signal Processing*, 2017.
- [9] Axel Berg, Mark O’Connor, and Miguel Tairum Cruz. Keyword transformer: A self-attention model for keyword spotting. *arXiv preprint arXiv:2104.00769*, 2021.
- [10] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *Proc. ICML*, 2021.
- [11] P Boyes Braem and RL Sutton-Spence. *The Hands Are The Head of The Mouth. The Mouth as Articulator in Sign Languages*. Hamburg: Signum Press, 2001. ISBN 3927731838.
- [12] Hannah Bull, Triantafyllos Afouras, Gül Varol, Samuel Albanie, Liliane Momeni, and Andrew Zisserman. Aligning subtitles in sign language videos. In *arxiv*, 2021.
- [13] Dogan Can and M. Saraçlar. Lattice indexing for spoken term detection. *IEEE Transactions on Audio, Speech, and Language Processing*, 19:2338–2347, 2011.

- [14] Jingyuan Chen, Xinpeng Chen, L. Ma, Zequn Jie, and Tat-Seng Chua. Temporally grounding natural sentence in video. In *EMNLP*, 2018.
- [15] Xi Chen, Shouyi Yin, Dandan Song, Peng Ouyang, Leibo Liu, and Shaojun Wei. Small-footprint keyword spotting with graph convolutional network. In *IEEE Automatic Speech Recognition and Understanding Workshop*, 2019.
- [16] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Proc. ACCV*, 2016.
- [17] Joon Son Chung and Andrew Zisserman. Signs in time: Encoding human motion as a temporal image. In *Workshop on Brave New Ideas for Motion Representations, ECCV*, 2016.
- [18] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. In *Proc. CVPR*, 2017.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805, 2018.
- [20] Runwei Ding, Cheng Pang, and Hong Liu. Audio-visual keyword spotting based on multidimensional convolutional neural network. In *IEEE International Conference on Image Processing*, 2018.
- [21] Linhao Dong, Shuang Xu, and Bo Xu. Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition. *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5884–5888, 2018.
- [22] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- [23] Santiago Fernandez, Alex Graves, and Jurgen Schmidhuber. An application of recurrent neural networks to discriminative keyword spotting. In *International Conference on Artificial Neural Networks*, 2007.
- [24] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *ICCV*, 2017.
- [25] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. Excl: Extractive clip localization using natural language descriptions. In *NAACL-HLT*, 2019.
- [26] John N Gowdy, Amarnag Subramanya, Chris Bartels, and Jeff Bilmes. Dbn based multi-stream models for audio-visual speech recognition. In *2004 IEEE International conference on acoustics, speech, and signal processing*, volume 1, pages I–993. IEEE, 2004.
- [27] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. Conformer: Convolution-augmented transformer for speech recognition. In *INTERSPEECH*, 2020.

- [28] Yanzhang He, Rohit Prabhavalkar, Kanishka Rao, Wei Li, Anton Bakhtin, and Ian McGraw. Streaming small-footprint keyword spotting using sequence-to-sequence models. In *Proceedings of IEEE ASRU*, 2017.
- [29] Lisa Anne Hendricks, O. Wang, E. Shechtman, Josef Sivic, Trevor Darrell, and Bryan C. Russell. Localizing moments in video with natural language. In *ICCV*, 2017.
- [30] Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, Jürgen Schmidhuber, et al. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [31] Kyuyeon Hwang, Minjae Lee, and Wonyong Sung. Online keyword spotting with a character-level recurrent neural network. *arXiv:1512.08903*, 2015.
- [32] Fumitada Itakura. *Minimum Prediction Residual Principle Applied to Speech Recognition*, page 154–158. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1990. ISBN 1558601244.
- [33] Abhishek Jha, Vinay P. Nambodiri, and C. V. Jawahar. Word spotting in silent lip videos. In *IEEE Winter Conference on Applications of Computer Vision*, 2018.
- [34] Prajwal K R, Triantafyllos Afouras, Andrew Zisserman, et al. Sub-word level lip reading with visual attention. *arXiv preprint arXiv:2110.07603*, 2021.
- [35] Shigeki Karita, Nelson Yalta, Shinji Watanabe, M. Delcroix, A. Ogawa, and T. Nakatani. Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration. In *INTERSPEECH*, 2019.
- [36] Taejun Kim and Juhan Nam. Temporal feedback convolutional recurrent neural networks for keyword spotting. *arXiv:1911.01803*, 2019.
- [37] Diederik P Kingma and Jimmy Ba. ADAM: A method for stochastic optimization. In *Proc. ICLR*, 2015.
- [38] Chris Lengerich and Awni Hannun. An end-to-end architecture for keyword spotting and voice activity detection. In *NIPS 2016 End-to-End Learning for Speech and Audio Processing Workshop*, 2016.
- [39] Wei Li, Sicheng Wang, Ming Lei, Sabato Marco Siniscalchi, and Chin-Hui Lee. Improving audio-visual speech recognition performance with cross-modal student-teacher training. In *Proc. ICASSP*, pages 6560–6564. IEEE, 2019.
- [40] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018.
- [41] Takaki Makino, Hank Liao, Yannis Assael, Brendan Shillingford, Basilio Garcia, Otavio Braga, and Olivier Siohan. Recurrent neural network transducer for audio-visual speech recognition. In *IEEE Workshop on Automatic Speech Recognition and Understanding*, 2019.
- [42] Abdelrahman Mohamed, Dmytro Okhonko, and Luke Zettlemoyer. Transformers with convolutional context for asr. *arXiv preprint arXiv:1904.11660*, 2019.

- [43] Liliane Momeni, Triantafyllos Afouras, Themis Stafylakis, Samuel Albanie, and Andrew Zisserman. Seeing wake words: Audio-visual keyword spotting. In *BMVC*, 2020.
- [44] Liliane Momeni, Gül Varol, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Watch, read and lookup: Learning to spot signs from multiple supervisors. In *Proc. ACCV*, 2020.
- [45] Dimitri Palaz, Gabriel Synnaeve, and Ronan Collobert. Jointly learning to locate and classify words using convolutional networks. In *INTERSPEECH*, 2016.
- [46] George Papandreou, Athanassios Katsamanis, Vassilis Pitsikalis, and Petros Maragos. Adaptive multimodal fusion by uncertainty compensation with application to audiovisual speech recognition. *Audio, Speech, and Language Processing, IEEE Transactions on*, 17(3):423–435, 2009.
- [47] Stavros Petridis, Themis Stafylakis, Pingchuan Ma, Georgios Tzimiropoulos, and Maja Pantic. Audio-Visual Speech Recognition with a Hybrid CTC/Attention Architecture. In *IEEE Spoken Language Technology Workshop*, pages 513–520, 2018.
- [48] Gerasimos Potamianos, Chalapathy Neti, Guillaume Gravier, Ashutosh Garg, and Andrew W Senior. Recent advances in the automatic recognition of audiovisual speech. *Proceedings of the IEEE*, 2003.
- [49] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [50] Richard C Rose and Douglas B Paul. A hidden markov model based keyword recognition system. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 129–132. IEEE, 1990.
- [51] Andrew Rosenberg, Kartik Audhkhasi, Abhinav Sethy, Bhuvana Ramabhadran, and Michael Picheny. End-to-end speech recognition and keyword search on low-resource languages. In *ICASSP*, 2017.
- [52] Tara N. Sainath and Carolina Parada. Convolutional neural networks for small-footprint keyword spotting. In *INTERSPEECH*, 2015.
- [53] Hiroaki Sakoe and Seibi Chiba. Dynamic programming algorithm optimization for spoken word recognition. *IEEE TRANSACTIONS ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING*, 26:43–49, 1978.
- [54] Adam Schembri, Jordan Fenlon, Ramas Rentelis, Sally Reynolds, and Kearsy Cormier. Building the British Sign Language Corpus. *Language Documentation & Conservation*, 7:136–154, 2013.
- [55] Adam Schembri, Jordan Fenlon, Ramas Rentelis, and Kearsy Cormier. British Sign Language Corpus Project: A corpus of digital video data and annotations of British Sign Language 2008-2017 (Third Edition), 2017. URL <http://www.bslcorpusproject.org>.

- [56] Brendan Shillingford, Yannis Assael, Matthew W. Hoffman, Thomas Paine, Cían Hughes, Utsav Prabhu, Hank Liao, Hasim Sak, Kanishka Rao, Lorraine Bennett, Marie Mulville, Ben Coppin, Ben Laurie, Andrew Senior, and Nando de Freitas. Large-Scale Visual Speech Recognition. In *INTERSPEECH*, 2019.
- [57] Speech Group at Carnegie Mellon University. CMU pronouncing dictionary. <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>, 2014.
- [58] Themis Stafylakis and Georgios Tzimiropoulos. Combining residual networks with lstms for lipreading. In *INTERSPEECH*, 2017.
- [59] Themis Stafylakis and Georgios Tzimiropoulos. Zero-shot keyword spotting for visual speech recognition in-the-wild. In *ECCV*, 2018.
- [60] Ming Sun, Anirudh Raju, George Tucker, Sankaran Panchapagesan, Gengshen Fu, Arindam Mandal, Spyridon Matsoukas, Nikko Strom, and Shiv Vitaladevuni. Max-pooling loss training of long short-term memory networks for small-footprint keyword spotting. *2016 IEEE Spoken Language Technology Workshop (SLT)*, 2016.
- [61] Rachel Sutton-Spence. Mouthings and simultaneity in british sign language. *Amsterdam studies in the theory and history of linguistic science, series 4*, 281:147, 2007.
- [62] Gül Varol, Liliame Momeni, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Read and attend: Temporal localisation in sign language videos. In *CVPR*, 2021.
- [63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [64] Yuxuan Wang, Pascal Getreuer, Thad Hughes, Richard Lyon, and Rif Saurous. Trainable frontend for robust and far-field keyword spotting. In *ICASSP*, 2017.
- [65] Jay Wilpon, Chin-Hui Lee, and Lawrence Rabiner. Application of hidden markov models for recognition of a limited set of words in unconstrained speech. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 254 – 257 vol.1, 06 1989. doi: 10.1109/ICASSP.1989.266413.
- [66] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Masayoshi Tomizuka, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision. *CoRR*, abs/2006.03677, 2020.
- [67] Pingping Wu, Hong Liu, Xiaofei Li, Ting Fan, and Xuewu Zhang. A novel lip descriptor for audio-visual keyword spotting based on adaptive decision fusion. *IEEE Transactions on Multimedia*, 2016.
- [68] Huijuan Xu, Kun He, Bryan A. Plummer, L. Sigal, S. Sclaroff, and Kate Saenko. Multilevel language and vision integration for text-to-clip retrieval. In *AAAI*, 2019.
- [69] Yue Yao, Tianyu Wang, Heming Du, Liang Zheng, and Tom Gedeon. Spotting visual keywords from temporal sliding windows. In *Mandarin Audio-Visual Speech Recognition Challenge*, 2019.

- [70] Jianwei Yu, Shi-Xiong Zhang, Jian Wu, Shahram Ghorbani, Bo Wu, Shiyin Kang, Shansong Liu, Xunying Liu, Helen Meng, and Dong Yu. Audio-visual recognition of overlapped speech for the lrs2 dataset. In *ICASSP*, pages 6984–6988, 05 2020. doi: 10.1109/ICASSP40776.2020.9054127.
- [71] Yitian Yuan, T. Mei, and Wenwu Zhu. To find where you talk: Temporal sentence localization in video with attention based location regression. In *AAAI*, 2019.
- [72] Runhao Zeng, H. Xu, W. Huang, Peihao Chen, Mingkui Tan, and Chuang Gan. Dense regression network for video grounding. In *CVPR*, 2020.
- [73] Haitong Zhang, Junbo Zhang, and Yujun Wang. Sequence-to-sequence models for small-footprint keyword spotting. *arXiv:1811.00348*, 2018.
- [74] Xingxuan Zhang, Feng Cheng, and Shilin Wang. Spatio-temporal fusion based convolutional sequence learning for lip reading. In *Proc. ICCV*, pages 713–722, 2019.
- [75] Yaodong Zhang and James Glass. Unsupervised spoken keyword spotting via segmental dtw on gaussian posteriorgrams. In *Proceedings of the 2009 IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 398 – 403, 01 2010. doi: 10.1109/ASRU.2009.5372931.
- [76] Yundong Zhang, Naveen Suda, Liangzhen Lai, and Vikas Chandra. Hello edge: Keyword spotting on microcontrollers. *CoRR*, 2017.
- [77] Ziheng Zhou, Guoying Zhao, Xiaopeng Hong, and Matti Pietikäinen. A review of recent advances in visual speech decoding. *Image and vision computing*, 32(9):590–605, 2014.
- [78] Yimeng Zhuang, Xuankai Chang, Yanmin Qian, and Kai Yu. Unrestricted Vocabulary Keyword Spotting Using LSTM-CTC. In *INTERSPEECH*, 2016.

A Application to silent films

Our Transpotter model has been trained to spot words in silent talking face videos. Naturally, the next step is to apply it on scenes from actual silent films to detect if a word is spoken and if so, also localize it in time. On the project website, we show qualitative examples of applying the Transpotter on scenes from the silent film, [The Artist](#). Silent films often use title cards to convey character dialogues. Although audio was not recorded, actors would still talk during the shooting, following the script, as it enabled them to act more naturally. We therefore used the film’s screenplay (available online) to identify scenes where title cards appear, which indicates the presence of dialogue, and queried words out of the title card content of these scenes. This application is particularly challenging because the title cards are only a guideline for the dialogue, therefore the lines that the actors end up using may differ substantially. Moreover we note that in this particular example, some of the actors are not native English speakers, therefore the model also has to deal with a domain shift in terms of accent. Despite those challenges, the Transpotter is still able to detect and properly localize words. Another application with a similar domain shift is spotting mouthings in sign language, which we have already discussed in the main paper.

The rest of this PDF is organized as follows. Section B describes the datasets we used for training and evaluation in detail. Next, we elaborate on the training process for the VSR baseline in Section C. In Section D, we describe the hyper-parameters choices for the Transpotter architecture. In Section E, we evaluate the model for longer query sequences containing multiple words. In Section F, we provide details for the architecture variants that were used in the main paper. In Section G, we conduct two more ablation studies, on the choice of the localization head and on the need for modality embeddings.

B Dataset Statistics

We conduct most of our experiments on the Lip Reading LRS2 [2, 17] and LRS3 [10] datasets. Each dataset contains three splits: (i) Pre-train, (ii) Train-val, and (iii) Test. The utterances in the pre-train set correspond to part-sentences as well as multiple sentences, whereas the training set only consists of single full sentences or phrases. The number of utterances, words and vocabulary for each of the splits in both datasets is reported in Table 5.

Dataset	Split	#Utterances	#Words	#Hours	Vocabulary
LRS2 [2, 17]	Pre-train	96k	2M	195	41k
	Train-val	47k	336k	29	18k
	Test	1.2k	6k	0.5	1.7k
LRS3 [10]	Pre-train	132k	3.9M	444	51k
	Train-val	32k	358k	30	17k
	Test	1.3k	10k	1	2k

Table 5: **Audio-visual datasets:** Statistics of LRS2 and LRS3 datasets.

Our models are first pre-trained on both the LRS2, LRS3 pre-train splits. Before testing on the test set of each dataset, we fine-tune the model on the train-val split of that particular

dataset. Training hyper-parameters are detailed in the next sections.

C VSR baseline

In this section, we describe the exact architectural details and training hyper-parameters for the VSR baseline.

Architecture details. The architecture closely resembles the TM-seq2seq architecture [9]. The only minor changes are all in the visual backbone, which we describe below in Table 6.

Layer	# channels	Kernel	Stride	Padding	Output dimensions
input	3	-	-	-	$T \times 112 \times 112$
conv _{1,1}	64	(5,5,5)	(1,2,2)	(2,2,2)	$T \times 56 \times 56$
conv _{2,1}	128	(3,3)	(2,2)	(1,1)	$T \times 28 \times 28$
conv _{2,2}	128	(3,3)	(1,1)	(1,1)	$T \times 28 \times 28$
conv _{2,3}	128	(3,3)	(1,1)	(1,1)	$T \times 28 \times 28$
conv _{3,1}	256	(3,3)	(2,2)	(1,1)	$T \times 14 \times 14$
conv _{3,2}	256	(3,3)	(1,1)	(1,1)	$T \times 14 \times 14$
conv _{3,3}	256	(3,3)	(1,1)	(1,1)	$T \times 14 \times 14$
conv _{4,1}	512	(3,3)	(2,2)	(1,1)	$T \times 7 \times 7$
conv _{4,2}	512	(3,3)	(1,1)	(1,1)	$T \times 7 \times 7$
conv _{4,3}	512	(3,3)	(1,1)	(1,1)	$T \times 7 \times 7$
conv _{5,1}	512	(3,3)	(2,2)	(1,1)	$T \times 4 \times 4$
fc	512	(4,4)	(1,1)	(0,0)	$T \times 1 \times 1$

Table 6: **Architecture details for the visual CNN backbone.** Batch Normalization and ReLU activation are added after every convolutional layer. Shortcut connections are also added at each layer, except for the first layer of every residual block – i.e. the ones with stride 2 (conv_{1,1}, conv_{2,1} and conv_{3,1}, conv_{4,1}).

Training hyper-parameters. We train the VSR baseline on LRS2, LRS3. The training consists of two stages, similar to [9].

In the first stage, we train the visual backbone end-to-end with the transformer layers on all two-word video clips (obtained using the available LRS2, LRS3 word alignments). Both the CNN and Transformer layers use the Adam optimizer [14], but with different learning rate schedules. For the CNN, we start with an initial learning rate of $1e^{-4}$ and reduce it by a factor of 2 every time the validation loss plateaus for 3 epochs. The minimum learning rate for the CNN is $1e^{-5}$. The transformer layers start with an initial learning rate of $5e^{-5}$. We do not reduce this learning rate in the first stage. The first stage takes approximately 6 days on four Tesla v100 GPUs.

After the visual backbone is trained, we extract the features for all the video clips in our datasets, and use these fixed features for all further training. We follow the curriculum strategy of [9] and train the transformer layers for longer video segments (9 words). We halve the learning rate every time the validation loss plateaus for 10 epochs. The minimum learning rate for the transformer layers is $1e^{-6}$. The second stage takes about 2 days on one Tesla v100 GPU.

Similar to the spotting models, we fine-tune on the train-val set of LRS2/LRS3 before evaluation. During inference, we decode text sequences with a left-to-right beam search with a beam width of $B = 20$. The beam hypotheses scores and textual outputs are used for computing the keyword spotting scores. More details of the evaluation mechanism can be found in this paper [28].

D Transpotter hyperparameters

Model hyper-parameters. We set the number of text encoder layers $N_t = 3$, the number of video encoder layers $N_v = 6$, and the number of joint multimodal transformer layers $T_m = 6$. The embedding dimension d is set to 512, and each multi-headed attention layer uses 8 heads.

Training hyper-parameters. The Transpotter is trained on the extracted visual features that we obtain after Stage 1 training (described above) of the VSR model. We use the Adam optimizer [37] with an initial learning rate of $5e^{-5}$, which is decreased by a factor of 5 every time the validation loss plateaus for 15 epochs. The minimum learning rate is $1e^{-6}$. The model is trained for 280 epochs on a single Ampere A40 GPU, which takes about 4 – 5 days. We fine-tune on the train-val set of LRS2/LRS3 before evaluation.

E Using the Transpotter to spot phrases

The Transpotter architecture can be used to spot multiple words (phrases) without any change in the architecture. Instead of feeding a phoneme sequence for a single word, we simply feed a phoneme sequence for a phrase. This is obtained by concatenating the phoneme sequence of each word in the phrase. We fine-tune the word-level model for a slight boost in performance.

In order to evaluate this model, we form our query set by using n-grams from the LRS2 and LRS3 test sets. We evaluate for $n = 1, 2, 3$. Note that $n = 1$ corresponds to single words. As is done for our single word evaluation in the main paper, we only use words with number of phonemes $n_p \geq 3$. The scores are reported in Table 7. We know that Transpotter becomes increasingly accurate with longer query lengths from Figure 2 (a) of the main paper, and the same trend continues when we try to spot phrases instead of single words. Spotting phrases is clearly far easier than spotting single words.

Num. words	LRS2				LRS3			
	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}
n = 1	65.0	87.1	69.2	68.3	52.0	77.1	55.4	53.6
n = 2	74.9	94.6	82.9	82.7	57.7	82.8	67.8	66.9
n = 3	87.9	99.4	92.9	92.7	74.7	91.0	82.0	80.6

Table 7: **Spotting phrases:** Transpotter can be trained to spot multi-word queries (phrases) with no change in the architecture. We report the scores for spotting unigrams (single words), bigrams and trigrams. Longer queries can be spotted far more easily.

F Details for the Architecture Ablations

In Table 2 (Section 4.3) of the main paper, we experiment with alternative Transformer architectures for this task. We explain each of the architectures below and also illustrate them

in Figure 6.

Enc_{vid}-Dec_{text}: This is the standard Transformer Encoder-Decoder architecture, which has also been used with great success [2] in related tasks such as VSR. The video encoder contains 6 Transformer layers encoding the temporal information for the video input. A text decoder consists of 6 Transformer decoder layers, each with (i) self-attention across the phoneme feature vectors, (ii) cross-attention between the phoneme and video feature vectors. The text input to the decoder is prepended with a [CLS] vector, where the probability of the keyword occurring in the video is predicted. Note that this model cannot explicitly localize the keyword in the video, as the output time-steps correspond to the text sequence and not the video.

Enc_{text}-Dec_{vid}: This architecture is similar to the previous one, except the inputs are swapped. The encoder processes the text sequence and the decoder inputs the video frames (prepended with a [CLS] vector). Since the output time-steps of the decoder correspond to the video, the exact same localization loss used in Transpotter can be applied here for localizing the keyword.

Transpotter w/o localization head: We also train our Transpotter model without the localization head. We only optimize for the presence/absence of the keyword at the [CLS] time-step.

All the above variants are trained using the exact same data pipeline and optimizer hyper-parameters.

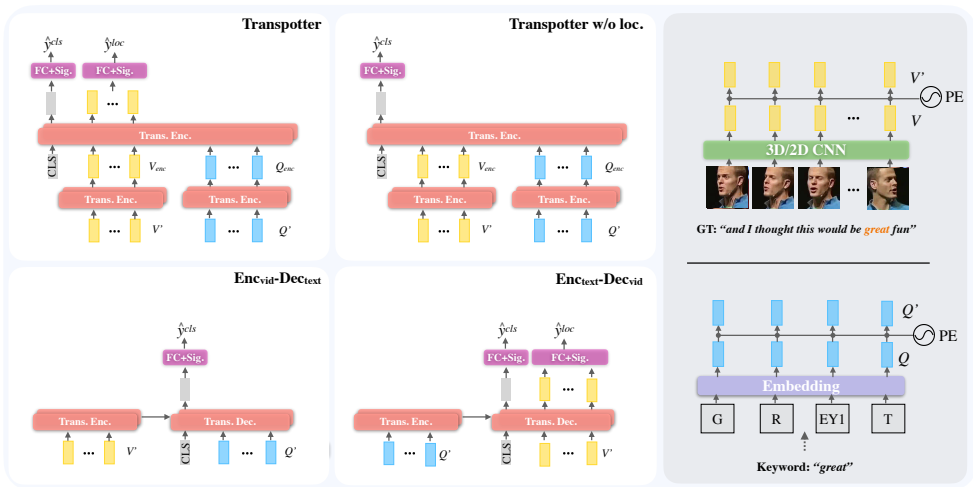


Figure 6: **Architecture ablations**: We illustrate each of the architecture variants tried for this task. The models are compared in Table 2 of the main paper.

G Additional Ablations

G.1 Choice of Localization Head

In Table G.1, we compare our proposed localization head with an alternative start-end span prediction head. For this, we have two MLP layers that predict start and end probabilities

at each video frame indicating the span of the keyword. The frame-wise probabilities are softmax-normalized and optimized with cross-entropy loss during training. We observe that our proposed localization head consistently outperforms this alternative technique.

Loss	LRS2				LRS3			
	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}
$L^{Cls} + L^{Loc}_{softmax}$	61.9	85.5	66.6	65.9	49.9	74.8	52.8	51.6
$L^{Cls} + L^{Loc}$	65.0	87.1	69.2	68.3	52.0	77.1	55.4	53.6

Table 8: **Ablations on the Localization head:** Per-frame sigmoid prediction of the keyword location outperforms the softmax-based span prediction.

G.2 Are modality embeddings needed?

Modality Tokens	LRS2				LRS3			
	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}	Acc _{@1} ^{Cls}	Acc _{@5} ^{Cls}	mAP ^{Cls}	mAP ^{Loc}
✓	64.6	87.3	68.9	68.3	52.9	77.0	55.4	53.5
×	65.0	87.1	69.2	68.3	52.0	77.1	55.4	53.6

Table 9: **Are modality embeddings needed in Transpotter?** Presence/absence of modality embeddings to explicitly specify the text and video time-steps does not affect the performance. Our dedicated video and text encoders encode sufficient modality specific information.

Since we concatenate our video and text embeddings together into a single sequence for feeding into the joint multi-modal transformer, the question arises if we need to explicitly specify which time-steps are video/text. We argue that this is not necessary because we have separate text and video encoders before the joint transformer. In order to empirically show this, we train a model with learnable modality tokens that are added to the text and video time-steps. We use two learnable vectors (one each for video/text) of embedding dimension d and add them to each of the video/text time-step accordingly. We perform this addition just before feeding them into the joint transformer. In Table G.2, we show that this performs comparably to our Transpotter model without any modality tokens.

H Error analysis on LRW

In this section, we illustrate some of the common errors made by our model. We conduct this analysis on the LRW test set using the same retrieval protocol that we have used in all our experiments, i.e. given a query word, we retrieve the top-K videos. Since LRW videos correspond to a single word label, we can clearly see which word labels are wrongly retrieved.

In Table H, we show the word labels of the incorrectly retrieved videos (among the top 10 retrievals) for some of the query words. We observe that often errors occur due to some of the phonemes of the query word appearing in the video. For example, the words “president” and “prison” share several phonemes. We also observe that compound words are a source of error: for example, in the second row, where “every” appears entirely in “everybody” and

“everything”. Finally, as discussed in Section 4.4, another failure case are homophemes. For example, in the last row, we see an example of the model retrieving the homopheme “million” for the query word “billion”.

Query word	Mistakes among the top 10 retrievals
president	press, prison, pressure
every	everybody, everything
allowed	allow, cloud, announced
example	couple, happen
billion	million, building

Table 10: **Error analysis on LRW:** For each query word, we report examples of incorrectly retrieved word instances among the top 10 retrievals. We can see clear patterns in mistakes such as phoneme overlap between the query word and the mistaken word, compound words and homophemes.