

Looking for the Signs: Identifying Isolated Sign Instances in Continuous Video Footage

Tao Jiang¹, Necati Cihan Camgöz¹, Richard Bowden¹

¹ Centre for Vision, Speech and Signal Processing,
University of Surrey, Guildford, UK

Abstract—In this paper, we focus on the task of one-shot sign spotting, i.e. given an example of an isolated sign (*query*), we want to identify whether/where this sign appears in a continuous, co-articulated sign language video (*target*). To achieve this goal, we propose a transformer-based network, called *Sign-Lookup*. We employ 3D Convolutional Neural Networks (CNNs) to extract spatio-temporal representations from video clips. To solve the temporal scale discrepancies between the query and the target videos, we construct multiple queries from a single video clip using different frame-level strides. Self-attention is applied across these query clips to simulate a continuous scale space. We also utilize another self-attention module on the target video to learn the contextual within the sequence. Finally a mutual-attention is used to match the temporal scales to localize the query within the target sequence. Extensive experiments demonstrate that the proposed approach can not only reliably identify isolated signs in continuous videos, regardless of the signers' appearance, but can also generalize to different sign languages. By taking advantage of the attention mechanism and the adaptive features, our model achieves state-of-the-art performance on the sign spotting task with accuracy as high as 96% on challenging benchmark datasets and significantly outperforming other approaches.

I. INTRODUCTION

Sign Languages are the native languages among Deaf communities. They are languages in their own right, distinct from spoken language but as linguistically complex as any spoken language. Each country has its own sign language often with regional variations. They incorporate manual (including handshape and motion) and non-manual (facial expression and body posture) channels or articulators which are combined via grammatical constructs that use both direction and space to convey meaning. As such, the grammar and word ordering of sign is very different to spoken language.

Although sign language dictionaries often give isolated forms of a sign, without any co-articulation or contextual modification, a continuous sign sequence is not just a concatenation of those isolated signs in spoken word order. Often, the first stage of linguistic annotation is to annotate this sequence of sign glosses contained in the video. Thus the ability to identify examples of signs within a continuous stream would be a valuable tool, both for linguistic analysis and computational sign language research.

In this work, we focus on the sign lookup problem, that is, looking for an isolated sign in a continuous, co-articulated sign language video. This process, often called *Sign Spotting*, has a wide range of applications e.g. searching in sign videos by keyword to allow a content-aware search;

creating or augmenting isolated sign dictionaries from other unlabeled footage; semi-automatic annotation (glossing) of sign language data for linguistic study; or used in assessment systems to automatically check whether an L2 Language learner is signing correctly.

One possible way to achieve this goal is to extract features from the *query* and the *target* videos, and then compute the feature similarity. However, directly applying this method will result in poor performance. This is because the *query* and the *target* videos may have severe spacial and temporal domains differences: (1) Signers in the *query* and *target* videos may be different people and/or have different appearance. (2) Different signers have their own signing habits or style, using slightly different motions even though they perform the same sign. (3) The sign speed in the *query* and the *target* videos may also be different. (4) The direction of the sign may be modified from the citation form of the sign (5) For one word, there might be multiple corresponding signing forms.

To tackle these challenges, we propose a Transformer-based deep learning model for sign lookup. We first utilize a 3D CNN to learn the spatio-temporal features from sign video clips. We pre-train this module for sign recognition using both the isolated signs (*query*) and continuous sign phrases (*target*). Once trained, the last classification layer is removed, and the 3D-CNN model is used as a feature extractor for the transformer model. To solve the temporal scale difference, we adopt different strategies while extracting features for the *query* and the *target* videos. For the isolated *query* sign, we extract 4 features with different temporal scales. Neighbouring frames in a window are separated from each other by a temporal offset of 1, 2, 4 or 8 frames. In this way, the extracted *query* features cover 4 different temporal scales, which we term adaptive features. For the continuous *target* video, a sliding window approach is applied to extract features throughout the whole video. Then self-attention is applied on the adaptive *query* features and the sliding *target* features. For the *query*, this simulates a continuous temporal scale space by learning a linear combination of discrete temporal scales. For the *target*, it learns the contextual relationship within the continuous sign video. In addition, mutual-attention learns to match multi-time scaled queries within the target sequence. This decreases the temporal scale difference between the *query* and the *target*. Finally, a multi-layer perceptron is use for classification.

The contributions of this paper can be listed as: (1) we propose a novel Transformer-based neural network for sign

lookup and demonstrate its application to sign spotting. (2) we extract adaptive features for a bank of isolated videos and apply self-attention, which simulates a continuous temporal scale space. (3) we apply mutual-attention on the adaptive *query* and *target* features, which closes the temporal domain difference gap between the *query* and the *target*. (4) We demonstrate generalization to two different sign languages and experiments demonstrate state-of-the-art performance.

The rest of the paper is structured as follows: In Section 2 we cover the related work on sign language recognition, sign spotting and sign language datasets. In Section 3 we describe the *Sign-Lookup* architecture and its components in detail. We describe the training process and implementation details in Section 4. In Section 5 we perform qualitative and quantitative experiments to demonstrate the effect of *Sign-Lookup* on sign language learning. We apply our model to the downstream task of sign spotting, achieving state-of-the-art performance. Finally, we conclude the paper in Section 6 by discussing our findings and the future work.

II. RELATED WORK

Our work is related to several topics in the literature including the curation of sign language datasets as well as sign language recognition and sign spotting. We will discuss each in turn.

Sign language datasets: There is no universal sign language. Each country has its own distinct language and may have regional dialects that have considerable difference even within a single country. However, some languages are better represented by public datasets than others. For example, there are many public datasets that have been released for American Sign Language (ASL) [2], [23], [31], [47], German Sign Language (DGS)[25], [45], Chinese Sign Language (CSL) [10], [22], Finnish Sign Language (FSL) [44] and British Sign Language (BSL) [1], [38] to name but a few. For BSL, Albanie *et al.* [1] use mouthing cues to collect a large-scale dataset of co-articulated BSL signs. However, this data is not currently publicly available. A large BSL dataset that is available and partially annotated is the BSLCORPUS captured by Schembri *et al.* [38]. The BSLCORPUS consists of continuous sign videos with fine-grained linguistic annotation. Felon *et al.* [14], [15] created an isolated BSL dictionary, for which the signs are consistent with the BSLCORPUS. As this provides a paired corpus of isolated signs and continuous video containing those signs, it is ideal for our application of sign spotting.

Sign language recognition: Automatic sign recognition has drawn a lot of attention from the computer vision research community in the last 30 years. Early work on sign language recognition [13], [16], [41], [24] focused on designing hand-crafted features that describe hand shape and motion. But to model the temporal variation of sign, approaches often mimic those developed for speech recognition. For example, Hidden Markov Models (HMMs) have been used widely [24], [40], [17], [46].

However, with the appearance of deep learning, neural networks have become popular. CNNs have demonstrated

significant improvement in manual [29], [28] and non-manual [26], [27] feature extraction. RNNs have shown superior performance to HMMs at learning dynamic temporal dependencies [32], [20], [49], but one problem for RNNs is they are prone to vanishing gradients. To overcome this problem, LSTM [48], [19], [35] and GRU [7] use gating to regulate the information flow. However, the information flow in an LSTM is one way: the current state can only inherit from its previous states; the following states have no effect on the current state. To overcome this, bidirectional LSTM (BLSTM) [6], [22], [50], [11] combine one forward LSTM and one backward LSTM such that the current state has full knowledge of the entire sequence. Just as 2D CNNs are ideal for image data, 3D CNNs [3], [5], [21], [49] have been applied to video sequences. When modelling spatio-temporal data, the depth dimension in 3D CNNs represents the temporal dimension of the sequence. I3D [9] based methods [1], [23] have shown significant performance improvements in sign language learning problems. In this paper, we use I3D to extract the spatio-temporal features from videos.

Sign language spotting: Viitaniemi *et al.* [44] crafted a sign descriptor based on a 5×5 skin distribution histogram, and used dynamic time warping for temporal alignment. Ong *et al.* [34] modeled the spatio-temporal signatures of a sign using Sequential Interval Patterns, and then organized them in a hierarchical tree structure for classification. According to the number of the instances available in the dictionary, sign spotting can be categorised as: zero-shot, one-shot or few-shot. For zero-shot learning, Bilge *et al.* [3] constructed spatio-temporal representations by using 3D CNNs + LSTMs, and then leveraged descriptive text embeddings to discover the unseen sign. Besides one-shot examples, work in [37] utilised a weakly supervised reservoir containing multiple instances of a gesture to significantly improve the performance of a gesture classifier. Momeni *et al.* [33] employed multiple-instance contrastive learning to solve the few-shot problem combined with other supervisory signals: sparse annotations and subtitles. In our case, where we have only a single *query* example of a sign, it is a one-shot learning problem.

III. SIGN-LOOKUP ARCHITECTURE

In this section, we introduce the architecture of *Sign-Lookup*, a deep learning network capable of identifying an isolated sign *query* within a continuous *target* sign video. Based on the extracted spatio-temporal features, our approach considers not only the temporal dependencies in a *target* video, but also mutual the dependencies between the *query* and *target* video.

The architecture of *Sign-Lookup* is shown in Figure 1. *Sign-Lookup* consists of three tiers of neural network. Firstly, a 3D Convolutional Neural Network (CNN) takes a sequence of images as input and extracts spatio-temporal features for both the *query* and *target* videos. Secondly, self-attention is applied to the *query* (adaptive) and the *target* features to exploit the temporal scale in the *query* video and the context in the *target* video. Mutual attention is employed to learn the dependencies between them. Finally, for classification, a

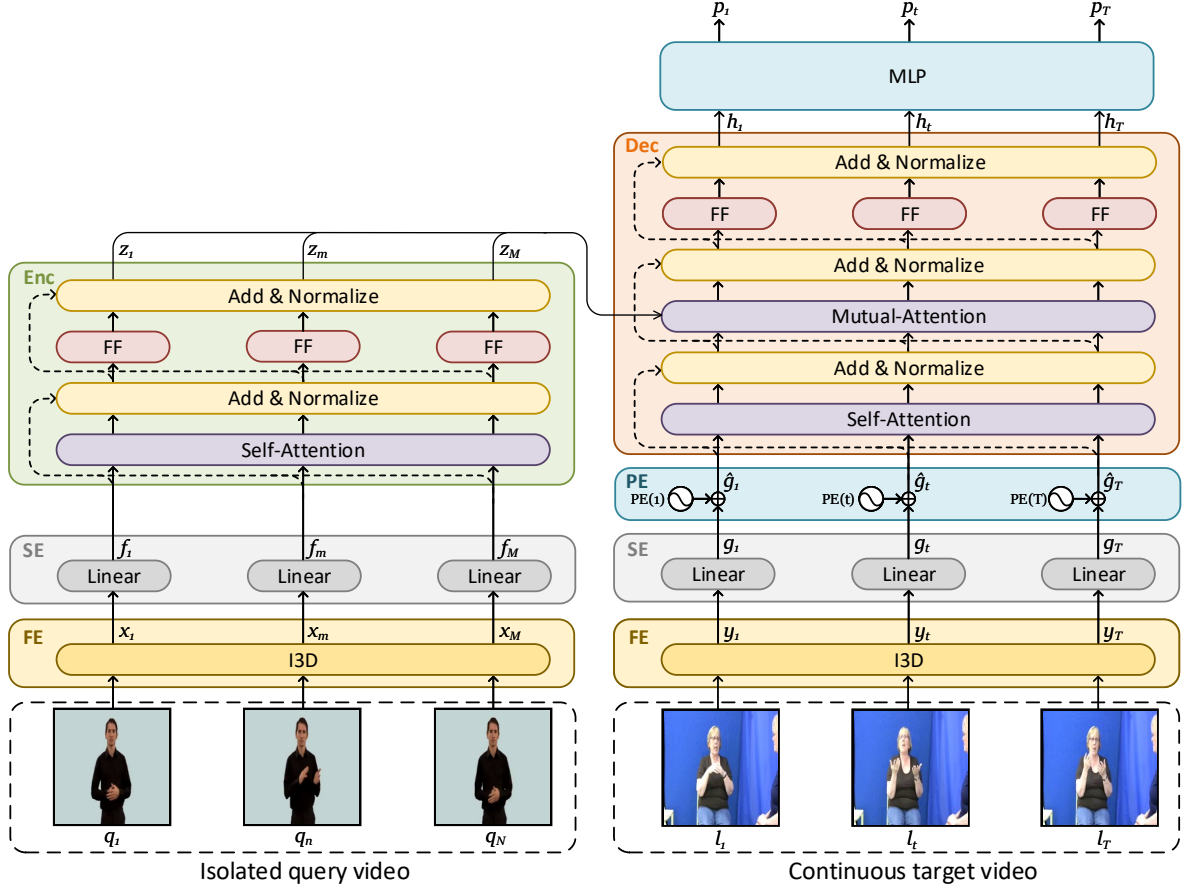


Fig. 1: A detailed overview of a single layered of Sign-Lookup.

(FE: Feature Extracting, SE: Spatial Embedding, PE: Positional Encoding, FF: Feed Forward, Enc: Encoder, Dec: Decoder)

multi-layer perceptron (MLP) is applied to make the final frame-level predictions. In the rest of this section, we give more details on each tier of *Sign-Lookup*.

A. Feature Extraction: 3D Convolutional Neural Networks

Sign-Lookup uses 3D CNNs to learn the spatio-temporal feature representations. Due to its successful application in human action recognition [12] and sign recognition [1], [23], we adopt the I3D architecture [9] to extract spatio-temporal features from the videos. Theoretically, the I3D model can take any length sequence of consecutive frames as input. However, according to [4], [36], [44], co-articulated signs roughly last for 13 frames, so, in our experiments, we set the window size of I3D to 16 frames. For a *query* video, we extract M adaptive features. For the m -th level feature, based on the middle frame of the video, we take a frame every $2^{(m-1)}$ frames bilaterally up to 16 frames. For the *target* video, we calculate a feature for each frame via a sliding window.

B. Learning Dependencies: Transformer

Sign spotting can be summarized as calculating the similarity between the features of the *query* and *target* sequence:

the more similar they are, the higher the chance the *query* is present in the *target*. However, we cannot just apply a dot product to the extracted features to calculate similarity, because: (1) The length of a signs vary and (2) The temporal scale is different between the *query* and the *target*.

To solve these problems, we propose a Transformer-based network. Assuming that we have a *query* video $Q = (q_1, q_2, \dots, q_N)$ with N frames and a *target* video $L = (l_1, l_2, \dots, l_T)$ with T frames, we extract M adaptive features $X = (x_1, x_2, \dots, x_M)$ for the *query* video and T features $Y = (y_1, y_2, \dots, y_T)$ for the *target* video via I3D (Fig1:FE). As can be seen in Figure 1 the features x_m and y_t are then projected into an embedding space through linear layers:

$$f_m = W^f \cdot x_m + b^f, \quad g_t = W^g \cdot y_t + b^g \quad (1)$$

To make the network aware of the frame ordering in the *target* sequence, we also use positional encoding [43] on the *target* embeddings:

$$\hat{g}_t = g_t + \text{PositionalEncoding}(t) \quad (2)$$

where PositionalEncoding is a predefined sinusoidal function conditioned on the sequence position t .

Next, the adaptive embedding vectors $F = (f_1, f_2, \dots, f_M)$ are fed into an encoder, where they are self-attended to learn a continuous temporal scale:

$$z_m = \text{Enc}(f_m | f_{1:M}) \quad (3)$$

where z_m denotes the contribution of the m -th basis f_m to the temporal scale space, which is generated by the encoder on the m -th time scale level, given the adaptive embeddings of the query video $f_{1:M}$.

The encoder consists of a stack of n identical layers. Each layer contains two sub-layers: Multi-Head Attention (MHA) and Feed-Forward (FF) sub-layer. MHA produces a weighted contextual representation, performing multiple projections of scaled dot-product attention. Scaled dot-product attention is a linear combination of values V , weighted by the relevant queries Q , keys K and dimensionality d_k :

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

The second Feed Forward sub-layer consists of two linear layers with ReLU activation in-between (Fig 1:FF).

The positional encoded vectors $\hat{G} = (\hat{g}_1, \hat{g}_2, \dots, \hat{g}_T)$ are fed into a decoder (Fig1:Dec). The decoder contains a self-attention module and a mutual-attention module. Self-attention learns the contextual relationships within the target sequence. Then mutual-attention learns the temporal scale relationship between the *query* and *target*. We formulate the decoding process as:

$$h_t = \text{Dec}(\hat{g}_t | \hat{g}_{1:T}, z_{1:M}) \quad (5)$$

where h_t denotes the projection of the t -th *target* feature in the temporal scale space of the *query*.

C. Classification: Multi-layer Perceptron

From these learned embeddings, we want a projection to the final classes which are dependant upon the final application. In the case of sign spotting, we have a binary decision indicating whether the *query* sign appears in the *target* video or not. To that end, we apply a multi-layer perceptron to the embedding. The MLP consists of 5 fully connected layers with Leaky Relu activation functions between them. The first linear layer has a residual connection on the 2048-dimensional embedding. The dimensionality of the embedding is then gradually reduced to 1024, 512, 256 and finally 2 for binary classification:

$$p_t = \text{MLP}(h_t) \quad (6)$$

where p_t denotes the probability of the query appearing in the t -th frame of the *target* video.

IV. TRAINING AND IMPLEMENTATION DETAILS

In this section, we discuss the implementation details necessary for training *Sign-Lookup*. First, we describe the sign video dataset used and then the pre-processing stages. As we use I3D for feature extraction, we next explain how to pre-train I3D before finally giving details of the transformer model parameters and its training process for sign spotting.

A. Sign Video Dataset

Our model makes use of two data streams: isolated sign videos which form the *query* and continuous *target* sign videos in which we seek to find the *query* sign. For the isolated sign videos, we use the publicly available BSL SignBank Dictionary introduced by [14], [15]. SignBank consists of 3688 lexical signs from BSL (i.e. signs that are highly conventionalised in both form and meaning across the sign language community). Each video in SignBank contains one prototypical canonical and isolated sign. However, most of the videos in SignBank are isolated examples, the signer remains still at both the start and the end lacking any form of co-articulation. As we do not want to include these useless frames during training we use OpenPose [8] to detect the signer’s hand movement and trim the beginning and the end of the video where the signer is silent. Then, the trimmed videos are used for both I3D pre-training and *Sign-Lookup* training. For the continuous videos, we use the ongoing BSL linguistic corpus BSLCORPUS introduced by [38]. The BSLCORPUS provides continuous BSL footage filmed by 249 Deaf participants from eight cities across the United Kingdom. It has partial human annotation for linguistic purposes. For sign spotting, we employ the gloss¹ annotation, which annotates approximate 72k signs with sign categories and temporal boundaries. For the annotated sign categories, there is correspondence between the SignBank lexicon and the gloss information annotated in the BSLCORPUS. This means we can use the isolated examples in SignBank to match against continuous examples in the BSLCORPUS. On this basis, we are only interested in the sign categories that are common to both datasets. Note that one English word might have different signing forms. Fortunately, SignBank and the BSLCORPUS distinguish these different forms consistently using different glosses. There are 2026 common sign categories for both datasets which we use in our tests. For pre-processing, we resize all videos in both datasets to 256 by 256 pixels with a fixed frame-rate of 25fps and then split the BSLCORPUS with a ratio of 4:1 for training and validation, respectively.

To demonstrate the generalization of *Sign-Lookup* across different sign languages, we also apply our model to DGS (German Sign Language). To do so we use the popular German Sign Language dataset: PHOENIX14T [7], which covers unconstrained DGS sign from 9 different signers with a vocabulary of 1066 different signs and translations into spoken German with a vocabulary of 2887 different words. It composes a parallel corpus including sign languages videos, sign-gloss annotation and German translation, which are all segmented into parallel sentences. For our purpose, we are only interested in the sign-gloss annotations.

B. I3D Pre-training

I3D was originally designed for action recognition so before using I3D as a feature extractor for sign, we need

¹A gloss is a written or typed approximation of a sign, typically using English words as “labels” for each sign and normally relating to a set lexicon.

to pre-train the I3D model on sign. Videos in SignBank are isolated and can be used for training directly. However, videos in the BSLCORPUS are continuous and cannot be used directly. To adapt them for I3D training, we segment the continuous videos into individual sign clips according to the annotated timeline for the corpus. We only do this for the training subset of the BSLCORPUS. Then, both SignBank training and the isolated BSLCORPUS videos are used for I3D pre-training. After pre-training, the last linear layer of I3D is removed and the rest of the parameters are frozen for feature extraction. In our case, the extracted feature vector has 1024 dimensions.

C. Sign-Lookup Training

When constructing the training dataset for *Sign-Lookup*, we take four clips of 16 frames with 1, 2, 4, 8 steps from SignBank videos and a clip of 64 frames as a sliding window from BSLCORPUS videos with a stride of 32 frames. To ensure the data is balanced, we ensure that for every positive sample added to training, we add a negative sample (i.e. a clip from the BSLCORPUS that does not contain the sign). The ground truth is *true* if the frame in the BSLCORPUS clip is annotated as the same sign as the SignBank example, *false* otherwise.

Once pre-trained, the I3D model is employed for feature extraction. In a SignBank video, we have four adaptive feature vectors. For the 64 frames in a BSLCORPUS clip, 64 feature vectors are extracted by applying I3D as a sliding window. We use Binary Cross-Entropy for classification which is minimized by Stochastic Gradient Descent (SGD) with mini-batches of size 8, and an initial learning rate of 10^{-2} with a plateau scheduler. This has a patience of 20 steps and a decrease factor of 0.1. All parameters in the network are initialized by xavier initialization [18]. Training was done on machines with 128GB RAM and a Nvidia Titan X GPU.

D. Evaluation Metrics

Following the works of [12], [30], we evaluate all methods using the metrics of frame-wise accuracy (Acc) and the F1 score at overlapping thresholds k denoted by $F1@k$. The frame-wise accuracy is the ratio of correctly predicted frames to total frames, which has been widely used in action detection and segmentation. It penalizes the prediction disorder but is insensitive to over-segmentation errors. To overcome this, we use $F1@k$ score proposed by [30]. It also allows minor temporal shifts, which can be caused by annotator variability. The F1 score is the harmonic mean of precision and recall: $F1 = 2 \frac{prec \cdot recall}{prec + recall}$. The intersection over union (IoU) overlap between the prediction and the ground truth is compared to the threshold k to decide a true or false positive, which is used for precision and recall calculation.

V. QUANTITATIVE EXPERIMENTS

In this section, we evaluate our model on different tasks and report quantitative results. We first apply our model to sign spotting with different ablation configurations to

demonstrate the importance of each component in our network. We then compare our model to prior work on sign-spotting task. Finally, we evaluate our model on other sign language datasets to show the generalization ability of *Sign-Lookup* amongst different sign languages.

A. Sign Spotting

We evaluate our model on sign spotting and investigate the effect of (1) the number of heads in an attention layer, (2) the number of encoder/decoder layers, (3) different batch sizes, (4) the adaptive features, (5) different levels of dropout. In the following experiments, we use 4 attention heads, 8 encoder/decoder layers. We train our models using a batch size of 4 and without any dropout unless specified otherwise.

Number of Attention Heads: Compared to classical dot product attention, multi-head attention allows the model to learn and attend to multiple representation subspaces. In this experiment, we evaluate the effect of the number of heads in an attention layer. We train four networks with different numbers of attention heads (1, 2, 4, 8).

TABLE I: Quantitative results for comparing different numbers of attention heads.

Heads Number	Acc	F1@{25,50,75}		
1	91.4	79.4	69.8	50.4
2	91.7	80.0	70.7	51.6
4	91.8	80.3	70.8	51.9
8	91.7	80.4	70.7	51.8

As can be seen in Table I, increasing the number of heads initially improves the performance. However, after 4 heads, the performance starts to decrease. We believe this is because the exploration in representation sub-spaces becomes saturated and introducing more heads after that point does not bring any performance enhancement but rather encourages over-fitting. We will use 4 attention heads in the remainder of our experiments.

Number of Encoder/Decoder Layers: An encoder can consist of a stack of n identical encoder layers. The lower layers contain more local features while the higher layers contain more global features. Theoretically, the more layers in the encoder, the more context is exploited in the embedding. However, too many can cause other problems like over-fitting and inefficiency. In this experiment, we train our model with different encoder layers to identify the trade-off and the optimal configuration. Encoders with 1, 2, 4, 8 layers are trained respectively. All results are shown in Table II.

Similar to attention heads, increasing the number of layers initially improves performance. However, after 4 layers, the performance starts to decrease. We believe this is because more layers give us more global features at the beginning, which provides more information to the network for classification. After a certain point, global feature learning becomes saturated and introducing more heads can only lead to over-fitting. We will use 4 encoder/decoder layers in the remainder of our experiments.

TABLE II: Quantitative results for comparing different number of encoder/decoder layers.

Encoder Layers	Acc	F1@{25,50,75}			
1	91.8	77.9	68.4	49.4	
2	91.8	80.1	70.6	51.6	
4	91.8	81.1	71.8	52.6	
8	91.8	80.3	70.8	51.9	

TABLE III: Quantitative results for comparing different batch sizes.

Batch Sizes	Acc	F1@{25,50,75}			
1	91.1	78.2	68.4	49.4	
2	91.8	81.0	71.7	52.5	
4	91.8	81.1	71.8	52.6	
8	92.0	80.8	71.9	52.8	
16	91.8	82.0	69.9	50.9	

Batch Size: Larger batch sizes tend to produce smoother gradients at the price of decreased convergence rate. On the contrary, smaller batch sizes can enhance the convergence rate but introduces more noise to training. However, according to the recent studies on the information theory behind deep learning [42], [39], the noise introduced by smaller batch size can help the network to represent the data more efficiently. In this experiment, we train our network with different batch sizes subject to the GPU memory limits. All the results are shown in Table 3.

From the table, we can see that the largest batch size does not achieve the best performance, neither does the smallest. Too large a batch size slows convergence which tends to converge to a local optima. While too small, introduces too much noise, which can diminish performance. In the remainder of our experiments, we will use a batch size of 8.

Adaptive features: In our model, we use adaptive features for the query video because of the temporal scale difference between the *query* and the *target* videos. To show the advantage of adaptive features, we compare to using only one feature for the *query* video: a single feature extracted from its middle 16 frames. With a single feature in the *query*, we need to change our network structure. Firstly removing the encoder attention as there are no other features to attend to and secondly to replace the mutual-attention with a dot and element-wise product. The results are shown in Table VI.

One obvious observation from this experiment is that our method with adaptive features outperforms the method with only one feature. This is because using only one feature assumes the temporal scale is the same in the *query* and the *target* videos, which is not true in this scenario. On the contrary, our method takes advantage of adaptive features to reduce the temporal scale difference, which achieves much better results.

Dropout: In this experiment, we train our model with dropout ranging from 0.0 to 0.5. All results are shown in Table V. As can be expected, the method with zero dropout

TABLE IV: Quantitative results comparing methods with one query feature and adaptive query features.

Methods	Acc	F1@{25,50,75}			
One feature	91.1	66.8	58.1	42.1	
Adaptive features	92.1	84.9	75.7	56.0	

TABLE V: Quantitative results for comparing different dropout.

Dropout	Acc	F1@{25,50,75}			
0.0	92.0	80.8	71.9	52.8	
0.1	91.7	83.0	73.4	53.8	
0.2	91.6	84.3	74.7	54.6	
0.3	92.1	84.9	75.7	56.0	
0.4	91.6	84.7	75.2	54.9	
0.5	91.5	83.6	73.6	53.5	

produces poor results with 0.3 providing optimal results.

B. Comparison to Prior work

In this experiment, we compare our model with the method introduced in [33], which uses a cosine function as the similarity measurement between the *query* and the *target* features. If the similarity value is above a predefined threshold, the *query* will be regarded as spotted. Their provided pretrained model is trained jointly on BSL-1K [1] and BSLDict datasets. However, neither the BSL-1K dataset nor the training code are accessible publicly. Hence, we apply their pretrained model on the BSLCORPUS using the common dictionary between BSLDICT and SignBank, while also noting the possibility of getting improved results if trained on the applied datasets. The threshold used in this experiment is 0.7. The comparison results are shown in Table 6. From the table, we can see that our method significantly outperforms Watch-Read-Lookup [33] on the sign spotting task.

C. Transfer to Other Sign Language

To evaluate the generalization of our model to different datasets, we apply *Sign-Lookup* to the German Sign Language dataset: PHOENIX14T. This dataset has frame-level annotation for glosses, however, it lacks the corresponding sign dictionary for use as the *query*. To overcome this, we use one instance of each gloss category from the training subset as the *query*. As there might be multiple instances for one category, we choose the example with a frame length closest to 16.

The performance of our model applied to the PHOENIX14T dataset is shown in Table VII. From

TABLE VI: Comparison against related approach to sign spotting on BSLCORPUS dataset.

Methods	Acc	F1@{25,50,75}			
Watch-Read-Lookup [33]	78.6	46.5	31.1	10.9	
Our method	92.1	85.6	76.4	56.3	

TABLE VII: Quantitative results for transferring our model to Germany Sign Language dataset: Phoenix2014T.

Dataset	Acc	F1@{25,50,75}		
Phoenix2014T - DEV	96.1	93.3	89.8	76.8
Phoenix2014T - TEST	95.6	93.6	89.1	75.7

the table, we can see that it outperforms the BSLCORPUS dataset tests. This is because we are using continuous videos from the same corpus for both *query* and *target* and this means that the spatial and temporal domain differences are minimal. However, these numbers are potentially worse than they should be as the gloss annotation on PHOENIX14T (as published) is automatically generated and therefore incorporates annotation noise.

VI. QUALITATIVE EXPERIMENTS

In this section, we report qualitative results for the *Sign-Lookup* model. We share snapshot examples of spotting a sign in Figure 2, with more examples provided in supplementary material. On the left hand side, is shown the isolated sign video we wish to use as the *query*. On the right side, it is the continuous *target* video we wish to identify the sign in. In the *target* video, the top bar (labeled gt) shows the ground truth with red depicting the start and stop of the sign in the larger continuous sequence. The bottom bar shows our prediction in green. The curve denotes the probability that the queried sign is spotted. In the upper-left corner of both videos, the corresponding gloss annotation is shown.

In the first example, we see that our model can accurately identify the *query* in the *target* video even though they come from two different datasets that have large spatial and temporal domain differences. In the second example, we successfully apply our model to a DGS Phoenix2014T example, demonstrating the generalization to other sign languages and datasets.

VII. CONCLUSION

In this paper, we proposed a novel transformer-based network called *Sign-Lookup*, which given a single isolated example of a sign, can identify or ‘spot’ occurrences of that sign in continuous real world signing footage. This not only solves the temporal domain difference between the *query* and the *target* videos, but also addresses the alignment between the sliding windows and sign boundaries. The experiments showed that *Sign-Lookup* improves the accuracy of sign spotting, especially when we have temporal scale discrepancy, achieving state-of-the-art performance. By also applying *Sign-Lookup* to a popular DGS dataset, we demonstrate its capability to generalise to other sign languages.

In the future, we would like to expand our approach to multi-channel streams including hands, body pose and facial expression and mouthings. We also plan to release this as a tool for linguists to help with annotation.



Fig. 2: Some qualitative sign spotting examples produced by our model. The results in the top two rows use the SignBank dataset for the *query* (left) and BSLCORPUS dataset for the *target* (right). The results in the bottom two rows use PHOENIX14T for both the *query* and *target*.

REFERENCES

- [1] S. Albanie, G. Varol, L. Momeni, T. Afouras, J. S. Chung, N. Fox, and A. Zisserman. Bsl-1k: Scaling up co-articulated sign language recognition using mouthing cues. *arXiv preprint arXiv:2007.12131*, 2020.
- [2] V. Athitsos, C. Neidle, S. Sclaroff, J. Nash, A. Stefan, Q. Yuan, and A. Thangali. The american sign language lexicon video dataset. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–8. IEEE, 2008.
- [3] Y. C. Bilge, N. Ikizler-Cinbis, and R. G. Cinbis. Zero-shot sign language recognition: Can textual data uncover sign languages? *arXiv preprint arXiv:1907.10292*, 2019.
- [4] P. Buehler, A. Zisserman, and M. Everingham. Learning sign language by watching tv (using weakly aligned subtitles). In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2961–2968. IEEE, 2009.
- [5] N. C. Camgoz, S. Hadfield, O. Koller, and R. Bowden. Using convolutional 3d neural networks for user-independent continuous gesture recognition. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 49–54. IEEE, 2016.
- [6] N. C. Camgoz, S. Hadfield, O. Koller, and R. Bowden. Subnets: End-to-end hand shape and continuous sign language recognition. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3075–3084. IEEE, 2017.
- [7] N. C. Camgoz, S. Hadfield, O. Koller, R. Bowden, and H. Ney. Neural

- Sign Language Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [8] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh. Openpose: realtime multi-person 2d pose estimation using part affinity fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):172–186, 2019.
 - [9] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
 - [10] X. Chai, H. Wang, and X. Chen. The devisign large vocabulary of chinese sign language database and baseline evaluations. *Technical report VIPL-TR-14-SLR-001. Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS*, 2014.
 - [11] R. Cui, H. Liu, and C. Zhang. Recurrent convolutional neural networks for continuous sign language recognition by staged optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7361–7369, 2017.
 - [12] Y. A. Farha and J. Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3575–3584, 2019.
 - [13] A. Farhadi, D. Forsyth, and R. White. Transfer learning in sign language. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
 - [14] J. Fenlon, K. Cormier, R. Rentelis, A. Schembri, K. Rowley, R. Adam, and B. Woll. Bsl signbank: A lexical database of british sign language. london: Deafness, cognition and language research centre, university college london, 2014.
 - [15] J. Fenlon, K. Cormier, and A. Schembri. Building bsl signbank: The lemma dilemma revisited. *International Journal of Lexicography*, 28(2):169–206, 2015.
 - [16] H. Fillbrandt, S. Akyol, and K.-F. Kraiss. Extraction of 3d hand shape and posture from image sequences for sign language recognition. In *2003 IEEE International SOI Conference. Proceedings (Cat. No. 03CH37443)*, pages 181–186. IEEE, 2003.
 - [17] J. Forster, C. Oberdörfer, O. Koller, and H. Ney. Modality combination techniques for continuous sign language recognition. In *Iberian conference on pattern recognition and image analysis*, pages 89–99. Springer, 2013.
 - [18] X. Glorot and Y. Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
 - [19] D. Guo, W. Zhou, H. Li, and M. Wang. Hierarchical lstm for sign language translation. In *AAAI*, 2018.
 - [20] P. Gupta, K. Kautz, et al. Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural networks. In *CVPR*, volume 1, page 3, 2016.
 - [21] J. Huang, W. Zhou, H. Li, and W. Li. Sign language recognition using 3d convolutional neural networks. In *2015 IEEE international conference on multimedia and expo (ICME)*, pages 1–6. IEEE, 2015.
 - [22] J. Huang, W. Zhou, Q. Zhang, H. Li, and W. Li. Video-based sign language recognition without temporal segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
 - [23] H. R. V. Joze and O. Koller. Ms-asl: A large-scale data set and benchmark for understanding american sign language. *arXiv preprint arXiv:1812.01053*, 2018.
 - [24] T. Kadir, R. Bowden, E.-J. Ong, and A. Zisserman. Minimal training, large lexicon, unconstrained sign language recognition. In *BMVC*, pages 1–10, 2004.
 - [25] O. Koller, J. Forster, and H. Ney. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141:108–125, 2015.
 - [26] O. Koller, H. Ney, and R. Bowden. Read my lips: Continuous signer independent weakly supervised viseme recognition. In *European Conference on Computer Vision*, pages 281–296. Springer, 2014.
 - [27] O. Koller, H. Ney, and R. Bowden. Deep learning of mouth shapes for sign language. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 85–91, 2015.
 - [28] O. Koller, H. Ney, and R. Bowden. Deep hand: How to train a cnn on 1 million hand images when your data is continuous and weakly labelled. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3793–3802, 2016.
 - [29] O. Koller, O. Zargaran, H. Ney, and R. Bowden. Deep sign: hybrid cnn-hmm for continuous sign language recognition. In *Proceedings of the British Machine Vision Conference 2016*, 2016.
 - [30] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017.
 - [31] D. Li, C. Rodriguez, X. Yu, and H. Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1459–1469, 2020.
 - [32] S. Mitra and T. Acharya. Gesture recognition: A survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37(3):311–324, 2007.
 - [33] L. Momeni, G. Varol, S. Albanie, T. Afouras, and A. Zisserman. Watch, read and lookup: learning to spot signs from multiple supervisors. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
 - [34] E.-J. Ong, O. Koller, N. Pugeault, and R. Bowden. Sign spotting using hierarchical sequential patterns with temporal intervals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1923–1930, 2014.
 - [35] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4594–4602, 2016.
 - [36] T. Pfister, J. Charles, and A. Zisserman. Large-scale learning of sign language by watching tv (using co-occurrences). In *BMVC*, 2013.
 - [37] T. Pfister, J. Charles, and A. Zisserman. Domain-adaptive discriminative one-shot learning of gestures. In *European Conference on Computer Vision*, pages 814–829. Springer, 2014.
 - [38] A. Schembri, J. Fenlon, R. Rentelis, S. Reynolds, and K. Cormier. Building the british sign language corpus. *Language Documentation & Conservation*, 7:136–154, 2013.
 - [39] R. Shwartz-Ziv and N. Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.
 - [40] T. E. Starner. Visual recognition of american sign language using hidden markov models. Technical report, Massachusetts Inst Of Tech Cambridge Dept Of Brain And Cognitive Sciences, 1995.
 - [41] S. Tamura and S. Kawasaki. Recognition of sign language motion images. *Pattern Recognition*, 21(4):343–353, 1988.
 - [42] N. Tishby and N. Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5. IEEE, 2015.
 - [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
 - [44] V. Viitaniemi, T. Jantunen, L. Savolainen, M. Karppa, and J. Laaksonen. S-pot—a benchmark in spotting signs within continuous signing. In *Proceedings of the 9th international conference on Language Resources and Evaluation (LREC 2014)*, ISBN 978-2-9517408-8-4. European Language Resources Association (LREC), 2014.
 - [45] U. Von Agris, M. Knorr, and K.-F. Kraiss. The significance of facial features for automatic sign language recognition. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 1–6. IEEE, 2008.
 - [46] U. Von Agris, J. Zieren, U. Canzler, B. Bauer, and K.-F. Kraiss. Recent developments in visual sign language recognition. *Universal Access in the Information Society*, 6(4):323–362, 2008.
 - [47] R. Wilbur and A. C. Kak. Purdue rvl-slll american sign language database. 2006.
 - [48] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville. Describing videos by exploiting temporal structure. In *Proceedings of the IEEE international conference on computer vision*, pages 4507–4515, 2015.
 - [49] Y. Ye, Y. Tian, M. Huenerfauth, and J. Liu. Recognizing american sign language gestures from within continuous videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2064–2073, 2018.
 - [50] H. Zhou, W. Zhou, Y. Zhou, and H. Li. Spatial-temporal multi-cue network for continuous sign language recognition. In *AAAI*, pages 13009–13016, 2020.