



Watch, Read and Lookup: Learning to Spot Signs from Multiple Supervisors

Liliane Momeni, Gül Varol^(✉), Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman

Visual Geometry Group, University of Oxford, Oxford, UK
{liliane,gul,albanie,afouras,az}@robots.ox.ac.uk

Abstract. The focus of this work is *sign spotting*—given a video of an isolated sign, our task is to identify *whether* and *where* it has been signed in a continuous, co-articulated sign language video. To achieve this sign spotting task, we train a model using multiple types of available supervision by: (1) *watching* existing sparsely labelled footage; (2) *reading* associated subtitles (readily available translations of the signed content) which provide additional *weak-supervision*; (3) *looking up* words (for which no co-articulated labelled examples are available) in visual sign language dictionaries to enable novel sign spotting. These three tasks are integrated into a unified learning framework using the principles of Noise Contrastive Estimation and Multiple Instance Learning. We validate the effectiveness of our approach on low-shot sign spotting benchmarks. In addition, we contribute a machine-readable British Sign Language (BSL) dictionary dataset of isolated signs, BSLDICT, to facilitate study of this task. The dataset, models and code are available at our project page (<https://www.robots.ox.ac.uk/~vgg/research/bsldict/>).

1 Introduction

The objective of this work is to develop a *sign spotting* model that can identify and localise instances of signs within sequences of continuous sign language. Sign languages represent the natural means of communication for deaf communities [1] and sign spotting has a broad range of practical applications. Examples include: indexing videos of signing content by keyword to enable content-based search; gathering diverse dictionaries of sign exemplars from unlabelled footage for linguistic study; automatic feedback for language students via an “auto-correct” tool (e.g. “did you mean this sign?”); making voice activated wake word devices accessible to deaf communities; and building sign language datasets by automatically labelling examples of signs.

The recent marriage of large-scale, labelled datasets with deep neural networks has produced considerable progress in audio [2, 3] and visual [4, 5] keyword

L. Momeni, G. Varol and S. Albanie—Equal contribution.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-69544-6_18) contains supplementary material, which is available to authorized users.




Fig. 1. We consider the task of *sign spotting* in co-articulated, continuous signing. Given a query dictionary video of an isolated sign (e.g. “apple”), we aim to identify *whether* and *where* it appears in videos of continuous signing. The wide domain gap between dictionary examples of *isolated* signs and target sequences of *continuous* signing makes the task extremely challenging.

spotting in *spoken languages*. However, a direct replication of these keyword spotting successes in sign language requires a commensurate quantity of labelled data (note that modern audiovisual spoken keyword spotting datasets contain millions of densely labelled examples [6, 7]). Large-scale corpora of continuous, co-articulated¹ signing from TV broadcast data have recently been built [8], but the labels accompanying this data are: (1) sparse, and (2) cover a limited vocabulary.

It might be thought that a sign language dictionary would offer a relatively straightforward solution to the sign spotting task, particularly to the problem of covering only a limited vocabulary in existing large-scale corpora. But, unfortunately, this is not the case due to the severe *domain differences* between dictionaries and continuous signing in the wild. The challenges are that sign language dictionaries typically: (i) consist of *isolated signs* which differ in appearance from the *co-articulated* sequences of continuous signs (for which we ultimately wish to perform spotting); and (ii) differ in speed (are performed more slowly) relative to co-articulated signing. Furthermore, (iii) dictionaries only possess a few examples of each sign (so learning must be *low shot*); and as one more challenge, (iv) there can be multiple signs corresponding to a single keyword, for example due to regional variations of the sign language [9]. We show through experiments in Sect. 4, that directly training a sign spotter for continuous signing on dictionary examples, obtained from an internet-sourced sign language dictionary, does indeed perform poorly.

To address these challenges, we propose a unified framework in which sign spotting embeddings are learned from the dictionary (to provide broad coverage of the lexicon) in combination with two additional sources of supervision. In aggregate, these multiple types of supervision include: (1) *watching* sign language and learning from existing sparse annotations; (2) exploiting weak-supervision by *reading* the subtitles that accompany the footage and extracting candidates for signs that we expect to be present; (3) *looking up* words (for which we do not have labelled examples) in a sign language dictionary (see Fig. 2 for an overview). The recent development of large-scale, subtitled corpora of continuous signing providing sparse annotations [8] allows us to study this problem setting directly. We formulate our approach as a Multiple Instance Learning problem

¹ *Co-articulation* refers to changes in the appearance of the current sign due to neighbouring signs.

in which positive samples may arise from any of the three sources and employ Noise Contrastive Estimation [10] to learn a domain-invariant (valid across both isolated and co-articulated signing) representation of signing content.

We make the following six contributions: (1) We provide a machine readable British Sign Language (BSL) dictionary dataset of isolated signs, BSLDICT, to facilitate study of the sign spotting task; (2) We propose a unified Multiple Instance Learning framework for learning sign embeddings suitable for spotting from three supervisory sources; (3) We validate the effectiveness of our approach on a co-articulated sign spotting benchmark for which only a small number (low-shot) of isolated signs are provided as labelled training examples, and (4) achieve state-of-the-art performance on the BSL-1K sign spotting benchmark [8] (closed vocabulary). We show qualitatively that the learned embeddings can be used to (5) automatically mine new signing examples, and (6) discover “faux amis” (false friends) between sign languages.

2 Related Work

Our work relates to several themes in the literature: *sign language recognition* (and more specifically *sign spotting*), *sign language datasets*, *multiple instance learning* and *low-shot action localization*. We discuss each of these themes next.

Sign Language Recognition. The study of automatic sign recognition has a rich history in the computer vision community stretching back over 30 years, with early methods developing carefully engineered features to model trajectories and shape [11–14]. A series of techniques then emerged which made effective use of hand and body pose cues through robust keypoint estimation encodings [15–18]. Sign language recognition also has been considered in the context of sequence prediction, with HMMs [11, 13, 19, 20], LSTMs [21–24], and Transformers [25] proving to be effective mechanisms for this task. Recently, convolutional neural networks have emerged as the dominant approach for appearance modelling [21], and in particular, action recognition models using spatio-temporal convolutions [26] have proven very well-suited for video-based sign recognition [8, 27, 28]. We adopt the I3D architecture [26] as a foundational building block in our studies.

Sign Language Spotting. The sign language spotting problem—in which the objective is to find performances of a sign (or sign sequence) in a longer sequence of signing—has been studied with Dynamic Time Warping and skin colour histograms [29] and with Hierarchical Sequential Patterns [30]. Different from our work which learns representations from multiple weak supervisory cues, these approaches consider a fully-supervised setting with a single source of supervision and use hand-crafted features to represent signs [31]. Our proposed use of a dictionary is also closely tied to *one-shot/few-shot learning*, in which the learner is assumed to have access to only a handful of annotated examples of the target category. One-shot dictionary learning was studied by [18] – different to their approach, we explicitly account for dialect variations in the dictionary (and validate the improvements brought by doing so in Sect. 4). Textual descriptions

from a dictionary of 250 signs were used to study zero-shot learning by [32] – we instead consider the practical setting in which a handful of video examples are available per-sign (and make this dictionary available). The use of dictionaries to locate signs in subtitled video also shares commonalities with *domain adaptation*, since our method must bridge differences between the dictionary and the target continuous signing distribution. A vast number of techniques have been proposed to tackle distribution shift, including several adversarial feature alignment methods that are specialised for the few-shot setting [33, 34]. In our work, we explore the domain-specific batch normalization (DSBN) method of [35], finding ultimately that simple batch normalization parameter re-initialization is most effective when jointly training on two domains after pre-training on the bigger domain. The concurrent work of [36] also seeks to align representation of isolated and continuous signs. However, our work differs from theirs in several key aspects: (1) rather than assuming access to a large-scale labelled dataset of isolated signs, we consider the setting in which only a handful of dictionary examples may be used to represent a word; (2) we develop a generalised Multiple Instance Learning framework which allows the learning of representations from weakly aligned subtitles whilst exploiting sparse labels and dictionaries (this integrates cues beyond the learning formulation in [36]); (3) we seek to label and improve performance on co-articulated signing (rather than improving recognition performance on isolated signing). Also related to our work, [18] uses a “reservoir” of weakly labelled sign footage to improve the performance of a sign classifier learned from a small number of examples. Different to [18], we propose a multi-instance learning formulation that explicitly accounts for signing variations that are present in the dictionary.

Sign Language Datasets. A number of sign language datasets have been proposed for studying Finnish [29], German [37, 38], American [27, 28, 39, 40] and Chinese [22, 41] sign recognition. For British Sign Language (BSL), [42] gathered a corpus labelled with sparse, but fine-grained linguistic annotations, and more recently [8] collected BSL-1K, a large-scale dataset of BSL signs that were obtained using a mouthing-based keyword spotting model. In this work, we contribute BSLDICT, a dictionary-style dataset that is complementary to the datasets of [8, 42] – it contains only a handful of instances of each sign, but achieves a comprehensive coverage of the BSL lexicon with a 9K vocabulary (vs a 1K vocabulary in [8]). As we show in the sequel, this dataset enables a number of sign spotting applications.

Multiple Instance Learning. Motivated by the readily available sign language footage that is accompanied by subtitles, a number of methods have been proposed for learning the association between signs and words that occur in the subtitle text [15, 18, 43, 44]. In this work, we adopt the framework of Multiple Instance Learning (MIL) [45] to tackle this problem, previously explored by [15, 46]. Our work differs from these works through the incorporation of a dictionary, and a principled mechanism for explicitly handling sign variants, to guide the learning process. Furthermore, we generalise the MIL framework so that it can learn to further exploit sparse labels. We also conduct experiments




Fig. 2. The proposed Watch, Read and Lookup framework trains sign spotting embeddings with three cues: (1) *watching* videos and learning from sparse annotation in the form of localised signs (lower-left); (2) *reading* subtitles to find candidate signs that may appear in the source footage (top); (3) *looking up* corresponding visual examples in a sign language dictionary and aligning the representation against the embedded source segment (lower-right).

at significantly greater scale to make use of the full potential of MIL, considering more than two orders of magnitude more weakly supervised data than [15, 46].

Low-Shot Action Localization. This theme investigates semantic video localization: given one or more query videos the objective is to localize the segment in an untrimmed video that corresponds semantically to the query video [47–49]. Semantic matching is too general for the sign-spotting considered in this paper. However, we build on the temporal ordering ideas explored in this theme.

3 Learning Sign Spotting Embeddings from Multiple Supervisors

In this section, we describe the task of *sign spotting* and the three forms of supervision we assume access to. Let \mathcal{X}_{Σ} denote the space of RGB video segments containing a frontal-facing individual communicating in sign language Σ and denote by $\mathcal{X}_{\Sigma}^{\text{single}}$ its restriction to the set of segments containing a single sign. Further, let \mathcal{T} denote the space of subtitle sentences and $\mathcal{V}_{\Sigma} = \{1, \dots, V\}$ denote the *vocabulary*—an index set corresponding to an enumeration of written words that are equivalent to signs that can be performed in Σ^2 .

Our objective, illustrated in Fig. 1, is to discover all occurrences of a given keyword in a collection of continuous signing sequences. To do so, we assume access to: (i) a subtitled collection of videos containing continuous signing, $\mathcal{S} = \{(x_i, s_i) : i \in \{1, \dots, I\}, x_i \in \mathcal{X}_{\Sigma}, s_i \in \mathcal{T}\}$; (ii) a sparse collection of temporal sub-segments of these videos that have been annotated with their corresponding word, $\mathcal{M} = \{(x_k, v_k) : k \in \{1, \dots, K\}, v_k \in \mathcal{V}_{\Sigma}, x_k \in$

² Sign language dictionaries provide a word-level or phrase-level correspondence (between sign language and spoken language) for many signs but no universally accepted *glossing* scheme exists for transcribing languages such as BSL [1].

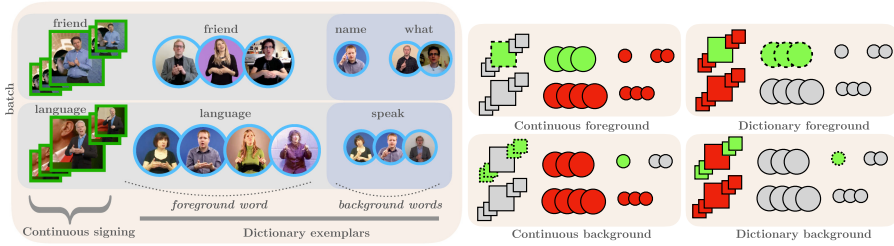


Fig. 3. Batch sampling and positive/negative pairs: We illustrate the formation of a batch when jointly training on continuous signing video (squares) and dictionaries of isolated signing (circles). **Left:** For each continuous video, we sample the dictionaries corresponding to the labeled word (foreground), as well as to the rest of the subtitles (background). **Right:** We construct positive/negative pairs by anchoring at 4 different portions of a batch item: continuous foreground/background and dictionary foreground/background. Positives and negatives (defined across continuous and dictionary domains) are green and red, respectively; anchors have a dashed border (see Appendix C.2 for details). (Color figure online)

$\mathcal{X}_{\Sigma}^{\text{single}}, \exists(x_i, s_i) \in \mathcal{S} \text{ s.t. } x_k \subseteq x_i$; (iii) a curated *dictionary* of signing instances $\mathcal{D} = \{(x_j, v_j) : j \in \{1, \dots, J\}, x_j \in \mathcal{X}_{\Sigma}^{\text{single}}, v_j \in \mathcal{V}_{\Sigma}\}$. To address the sign spotting task, we propose to learn a *data representation* $f : \mathcal{X}_{\Sigma} \rightarrow \mathbb{R}^d$ that maps video segments to vectors such that they are *discriminative* for sign spotting and *invariant* to other factors of variation. Formally, for any labelled pair of video segments $(x, v), (x', v')$ with $x, x' \in \mathcal{X}_{\Sigma}$ and $v, v' \in \mathcal{V}_{\Sigma}$, we seek a data representation, f , that satisfies the constraint $\delta_{f(x)f(x')} = \delta_{vv'}$, where δ represents the Kronecker delta.

3.1 Integrating Cues Through Multiple Instance Learning

To learn f , we must address several challenges. First, as noted in Sect. 1, there may be a considerable distribution shift between the dictionary videos of isolated signs in \mathcal{D} and the co-articulated signing videos in \mathcal{S} . Second, sign languages often contain multiple sign variants for a single written word (resulting from regional dialects and synonyms). Third, since the subtitles in \mathcal{S} are only weakly aligned with the sign sequence, we must learn to associate signs and words from a noisy signal that lacks temporal localisation. Fourth, the localised annotations provided by \mathcal{M} are sparse, and therefore we must make good use of the remaining segments of subtitled videos in \mathcal{S} if we are to learn an effective representation.

Given full supervision, we could simply adopt a pairwise metric learning approach to align segments from the videos in \mathcal{S} with dictionary videos from \mathcal{D} by requiring that f maps a pair of isolated and co-articulated signing segments to the same point in the embedding space if they correspond to the same sign (*positive* pairs) and apart if they do not (*negative* pairs). As noted above, in practice we do not have access to positive pairs because: (1) for any annotated segment $(x_k, v_k) \in \mathcal{M}$, we have a set of potential sign variations represented in the dictionary (annotated with the common label v_k), rather than a single

unique sign; (2) since \mathcal{S} provides only weak supervision, even when a word is mentioned in the subtitles we do not know where it appears in the continuous signing sequence (if it appears at all). These ambiguities motivate a Multiple Instance Learning [45] (MIL) objective. Rather than forming positive and negative pairs, we instead form positive *bags* of pairs, $\mathcal{P}^{\text{bags}}$, in which we expect at least one pairing between a segment from a video in \mathcal{S} and a dictionary video from \mathcal{D} to contain the same sign, and negative bags of pairs, $\mathcal{N}^{\text{bags}}$, in which we expect no (video segment, dictionary video) pair to contain the same sign. To incorporate the available sources of supervision into this formulation, we consider two categories of positive and negative bag formations, described next (due to space constraints, a formal mathematical description of the positive and negative bags described below is deferred to Appendix C.2).

Watch and Lookup: Using Sparse Annotations and Dictionaries. Here, we describe a baseline where we assume no subtitles are available. To learn f from \mathcal{M} and \mathcal{D} , we define each positive bag as the set of possible pairs between a *labelled (foreground)* temporal segment of a continuous video from \mathcal{M} and the examples of the corresponding sign in the dictionary (green regions in Fig. A.2). The key assumption here is that each labelled sign segment from \mathcal{M} matches *at least one* sign variation in the dictionary. Negative bags are constructed by (i) anchoring on a continuous foreground segment and selecting dictionary examples corresponding to different words from other batch items; (ii) anchoring on a dictionary foreground set and selecting continuous foreground segments from other batch items (red regions in Fig. A.2). To maximize the number of negatives within one minibatch, we sample a different word per batch item.

Watch, Read and Lookup: Using Sparse Annotations, Subtitles and Dictionaries. Using just the labelled sign segments from \mathcal{M} to construct bags has a significant limitation: f is not encouraged to represent signs beyond the initial vocabulary represented in \mathcal{M} . We therefore look at the subtitles (which contain words beyond \mathcal{M}) to construct additional bags. We determine more positive bags between the set of *unlabelled (background)* segments in the continuous footage and the set of dictionaries corresponding to the background words in the subtitle (green regions in Fig. 3, right-bottom). Negatives (red regions in Fig. 3) are formed as the complements to these sets by (i) pairing continuous background segments with dictionary samples that can be excluded as matches (through subtitles) and (ii) pairing background dictionary entries with the foreground continuous segment. In both cases, we also define negatives from other batch items by selecting pairs where the word(s) have no overlap, e.g., in Fig. 3, the dictionary examples for the background word ‘speak’ from the second batch item are negatives for the background continuous segments from the first batch item, corresponding to the unlabelled words ‘name’ and ‘what’ in the subtitle.

To assess the similarity of two embedded video segments, we employ a similarity function $\psi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ whose value increases as its arguments become more similar (in this work, we use cosine similarity). For notational convenience below, we write ψ_{ij} as shorthand for $\psi(f(x_i), f(x_j))$. To learn f , we consider a generalization of the InfoNCE loss [50, 51] (a non-parametric softmax loss for-

mulation of Noise Contrastive Estimation [10]) recently proposed by [52]:

$$\mathcal{L}_{\text{MIL-NCE}} = -\mathbb{E}_i \left[\log \frac{\sum_{(j,k) \in \mathcal{P}(i)} \exp(\psi_{jk}/\tau)}{\sum_{(j,k) \in \mathcal{P}(i)} \exp(\psi_{jk}/\tau) + \sum_{(l,m) \in \mathcal{N}(i)} \exp(\psi_{lm}/\tau)} \right], \quad (1)$$

where $\mathcal{P}(i) \in \mathcal{P}^{\text{bags}}$, $\mathcal{N}(i) \in \mathcal{N}^{\text{bags}}$, τ , often referred to as the *temperature*, is set as a hyperparameter (we explore the effect of its value in Sect. 4).

3.2 Implementation Details

In this section, we provide details for the learning framework covering the embedding architecture, sampling protocol and optimization procedure.

Embedding Architecture. The architecture comprises an I3D spatio-temporal trunk network [26] to which we attach an MLP consisting of three linear layers separated by leaky ReLU activations (with negative slope 0.2) and a skip connection. The trunk network takes as input 16 frames from a 224×224 resolution video clip and produces 1024-dimensional embeddings which are then projected to 256-dimensional sign spotting embeddings by the MLP. More details about the embedding architecture can be found in Appendix C.1.

Joint Pretraining. The I3D trunk parameters are initialised by pretraining for sign classification jointly over the sparse annotations \mathcal{M} of a continuous signing dataset (BSL-1K [8]) and examples from a sign dictionary dataset (BSLDICT) which fall within their common vocabulary. Since we find that dictionary videos of isolated signs tend to be performed more slowly, we uniformly sample 16 frames from each dictionary video with a random shift and random frame rate n times, where n is proportional to the length of the video, and pass these clips through the I3D trunk then average the resulting vectors before they are processed by the MLP to produce the final dictionary embeddings. We find that this form of random sampling performs better than sampling 16 consecutive frames from the isolated signing videos (see Appendix C.1 for more details). During pretraining, minibatches of size 4 are used; and colour, scale and horizontal flip augmentations are applied to the input video, following the procedure described in [8]. The trunk parameters are then frozen and the MLP outputs are used as embeddings. Both datasets are described in detail in Sect. 4.1.

Minibatch Sampling. To train the MLP given the pretrained I3D features, we sample data by first iterating over the set of labelled segments comprising the sparse annotations, \mathcal{M} , that accompany the dataset of continuous, subtitled sampling to form minibatches. For each continuous video, we sample 16 consecutive frames around the annotated timestamp (more precisely a random offset within 20 frames before, 5 frames after, following the timing study in [8]). We randomly sample 10 additional 16-frame clips from this video outside of the labelled window, i.e., continuous background segments. For each subtitled sequence, we sample the dictionary entries for all subtitle words that appear in \mathcal{V}_{Σ} (see Fig. 3 for a sample batch formation).

Table 1. Datasets: We provide (i) the number of individual sign videos, (ii) the vocabulary size of the annotated signs, and (iii) the number of signers for BSL-1K and BSLDICT. BSL-1K is large in the number of annotated signs whereas BSLDICT is large in the vocabulary size. Note that we use a different partition of BSL-1K with longer sequences around the annotations as described in Sect. 4.1.

Dataset	#Videos	Vocab	#Signers
BSL-1K[8]	273K	1,064	40
BSLDICT	14,210	9,283	>28

Our minibatch comprises 128 sequences of continuous signing and their corresponding dictionary entries (we investigate the impact of batch size in Sect. 4.3). The embeddings are then trained by minimising the loss defined in Eq. (1) in conjunction with positive bags, $\mathcal{P}^{\text{bags}}$, and negative bags, $\mathcal{N}^{\text{bags}}$, which are constructed on-the-fly for each minibatch (see Fig. 3).

Optimization. We use a SGD optimizer with an initial learning rate of 10^{-2} to train the embedding architecture. The learning rate is decayed twice by a factor of 10 (at epoch 40 and 45). We train all models, including baselines and ablation studies, for 50 epochs at which point we find that learning has always converged.

Test Time. To perform spotting, we obtain the embeddings learned with the MLP. For the dictionary, we have a single embedding averaged over the video. Continuous video embeddings are obtained with sliding window (stride 1) on the entire sequence. We calculate the cosine similarity score between the continuous signing sequence embeddings and the embedding for a given dictionary video. We determine the location with the maximum similarity as the location of the queried sign. We maintain embedding sets of all variants of dictionary videos for a given word and choose the best match as the one with the highest similarity.

4 Experiments

In this section, we first present the datasets used in this work (including the contributed BSLDICT dataset) in Sect. 4.1, followed by the evaluation protocol in Sect. 4.2. We illustrate the benefits of the *Watch, Read and Lookup* learning framework for sign spotting against several baselines with a comprehensive ablation study that validates our design choices (Sect. 4.3). Finally, we investigate three applications of our method in Sect. 4.4, showing that it can be used to (i) not only spot signs, but also identify the specific sign variant that was used, (ii) label sign instances in continuous signing footage given the associated subtitles, and (iii) discover “faux amis” between different sign languages.

4.1 Datasets

Although our method is conceptually applicable to a number of sign languages, in this work we focus primarily on BSL, the sign language of British deaf communities. We use BSL-1K [8], a large-scale, subtitled and sparsely annotated dataset

of more than 1000 h of continuous signing which offers an ideal setting in which to evaluate the effectiveness of the *Watch, Read and Lookup* sign spotting framework. To provide dictionary data for the *lookup* component of our approach, we also contribute BSLDICT, a diverse visual dictionary of signs. These two datasets are summarised in Table 1 and described in more detail below.

BSL-1K. [8] comprises a vocabulary of 1,064 signs which are sparsely annotated over 1,000 h of video of continuous sign language. The videos are accompanied by subtitles. The dataset consists of 273K localised sign annotations, automatically generated from sign-language-interpreted BBC television broadcasts, by leveraging weakly aligned subtitles and applying keyword spotting to signer *mouthings*. Please refer to [8] for more details on the automatic annotation pipeline. In this work, we process this data to extract long videos with subtitles. In particular, we pad $+/-2$ s around the subtitle timestamps and we add the corresponding video to our training set if there is a sparse annotation word falling within this time window, assuming that the signing is reasonably well-aligned with its subtitles in these cases. We further consider only the videos whose subtitle duration is longer than 2 s. For testing, we use the automatic test set (corresponding to mouthing locations with confidences above 0.9). Thus we obtain 78K training and 3K test videos, each of which has a subtitle of 8 words on average and 1 sparse mouthing annotation.

BSLDICT. BSL dictionary videos are collected from a BSL sign aggregation platform signbsl.com [53], giving us a total of 14,210 video clips for a vocabulary of 9,283 signs. Each sign is typically performed several times by different signers, often in different ways. The dictionary consists of at least 28 different signers: the videos are downloaded from 28 known website sources and each source has at least 1 signer. The dictionary videos are of isolated signs (as opposed to co-articulated in BSL-1K): this means (i) the start and end of the video clips usually consist of a still signer pausing, and (ii) the sign is performed at a much slower rate for clarity. We first trim the sign dictionary videos, using body keypoints estimated with OpenPose [54] which indicate the start and end of wrist motion, to discard frames where the signer is still. With this process, the average number of frames per video drops from 78 to 56 (still significantly larger than co-articulated signs). To the best of our knowledge, BSLDICT is the first curated, BSL sign dictionary dataset for computer vision research, which will be made available. For the experiments in which BSLDICT is filtered to the 1,064 vocabulary of BSL-1K (see below), we have a total of 2,992 videos. Within this subset, each sign has between 1 and 10 examples (average of 3).

4.2 Evaluation Protocols

Protocols. We define two settings: (i) training with the entire 1064 vocabulary of annotations in BSL-1K; and (ii) training on a subset with 800 signs. The latter is needed to assess the performance on novel signs, for which we do not have access to co-articulated labels at training. We thus use the remaining 264

Table 2. The effect of the loss formulation: Embeddings learned with the classification loss are suboptimal since they are not trained for matching the two domains. Contrastive-based loss formulations (NCE) significantly improve, particularly when we adopt the multiple-instance variant introduced as our Watch-Read-Lookup framework of multiple supervisory signals.

Embedding arch	Supervision	Train (1064)		Train (800)	
		Seen (264)		Unseen (264)	
		mAP	R@5	mAP	R@5
$I3D^{BSLDICT}$	Classification	2.68	3.57	1.21	1.29
$I3D^{BSL-1K}$ [8]	Classification	13.09	17.25	6.74	8.94
$I3D^{BSL-1K,BSLDICT}$	Classification	19.81	25.57	4.81	6.89
$I3D^{BSL-1K,BSLDICT}+MLP$	Classification	36.75	40.15	10.28	14.19
$I3D^{BSL-1K,BSLDICT}+MLP$	InfoNCE	42.52	53.54	10.88	14.23
$I3D^{BSL-1K,BSLDICT}+MLP$	Watch-Lookup	43.65	53.03	11.05	14.62
$I3D^{BSL-1K,BSLDICT}+MLP$	Watch-Read-Lookup	48.11	58.71	13.69	17.79

words for testing. This test set is therefore common to both training settings, it is either ‘seen’ or ‘unseen’ at training. However, we do not limit the vocabulary of the dictionary as a practical assumption, for which we show benefits.

Metrics. The performance is evaluated based on ranking metrics. For every sign s_i in the test vocabulary, we first select the BSL-1K test set clips which have a mouthing annotation of s_i and then record the percentage of dictionary clips of s_i that appear in the first 5 retrieved results, this is the ‘Recall at 5’ (R@5). This is motivated by the fact that different English words can correspond to the same sign, and vice versa. We also report mean average precision (mAP). For each video pair, the match is considered correct if (i) the dictionary clip corresponds to s_i and the BSL-1K video clip has a mouthing annotation of s_i , and (ii) if the predicted location of the sign in the BSL-1K video clip, i.e. the time frame where the maximum similarity occurs, lies within certain frames around the ground truth mouthing timing. In particular, we determine the correct interval to be defined between 20 frames before and 5 frames after the labeled time (based on the study in [8]). Finally, because BSL-1K test is class-unbalanced, we report performances averaged over the test classes.

4.3 Ablation Study

In this section, we evaluate different components of our approach. We first compare our contrastive learning approach with classification baselines. Then, we investigate the effect of our multiple-instance loss formulation. We provide ablations for the hyperparameters, such as the batch size and the temperature, and report performance on a sign spotting benchmark.

Table 3. Extending the dictionary vocabulary: We show the benefits of sampling dictionary videos outside of the sparse annotations, using subtitles. Extending the lookup to the dictionary from the subtitles to the full vocabulary of BSLDICT brings significant improvements for novel signs (the training uses sparse annotations for the 800 words, and the remaining 264 for test).

Supervision	Dictionary Vocab	mAP	R@5
Watch-Read-Lookup	800 training vocab	13.69	17.79
Watch-Read-Lookup	9k full vocab	15.39	20.87

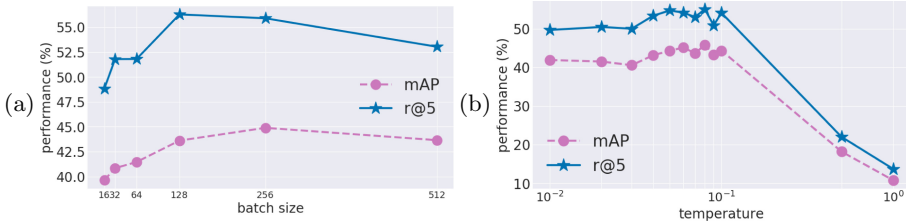


Fig. 4. The effect of (a) the **batch size** that determines the number of negatives across sign classes and (b) the **temperature** hyper-parameter for the MIL-NCE loss in Watch-Lookup against mAP and R@5 (trained on the full 1064 vocab.)

I3D Baselines. We start by evaluating baseline I3D models trained with classification on the task of spotting, using the embeddings before the classification layer. We have three variants in Table 2: (i) $I3D^{BSL-1K}$ provided by [8] which is trained only on the BSL-1K dataset, and we also train (ii) $I3D^{BSLDICT}$ and (iii) $I3D^{BSL-1K,BSLDICT}$. Training only on BSLDICT ($I3D^{BSLDICT}$) performs significantly worse due to the few examples available per class and the domain gap that must be bridged to spot co-articulated signs, suggesting that dictionary samples alone do not suffice to solve the task. We observe improvements with fine-tuning $I3D^{BSL-1K}$ jointly on the two datasets ($I3D^{BSL-1K,BSLDICT}$), which becomes our base feature extractor for the remaining experiments to train a shallow MLP.

Loss Formulation. We first train the MLP parameters on top of the frozen I3D trunk with classification to establish a baseline in a comparable setup. Note that, this shallow architecture can be trained with larger batches than I3D. Next, we investigate variants of our loss to learn a joint sign embedding between BSL-1K and BSLDICT video domains: (i) standard single-instance InfoNCE [50, 51] loss which pairs each BSL-1K video clip with *one* positive BSLDICT clip of the same sign, (ii) Watch-Lookup which considers multiple positive dictionary candidates, but does not consider subtitles (therefore limited to the annotated video clips). Table 2 summarizes the results. Our Watch-Read-Lookup formulation which effectively combines multiple sources of supervision in a multiple-instance framework outperforms the other baselines in both *seen* and *unseen* protocols.

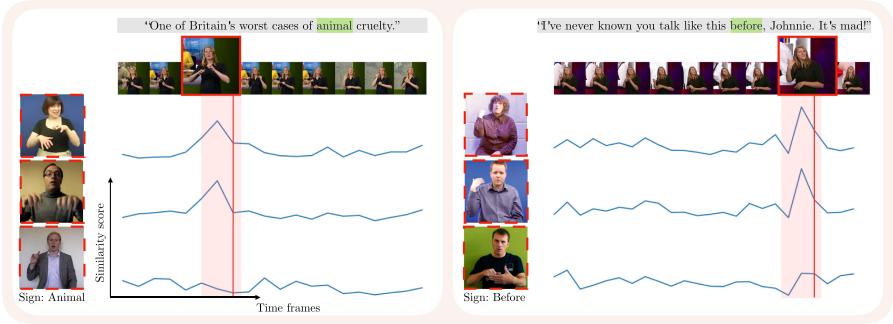


Fig. 5. Sign variant identification: We plot the similarity scores between BSL-1K test clips and BSLDICT variants of the sign “animal” (left) and “before” (right) over time. The labeled mouthing times are shown by red vertical lines and the sign proposal regions are shaded. A high similarity occurs for the first two rows, where the BSLDICT examples match the variant used in BSL-1K. (Color figure online)

Extending the Vocabulary. The results presented so far were using the same vocabulary for both continuous and dictionary datasets. In reality, one can assume access to the entire vocabulary in the dictionary, but obtaining annotations for the continuous videos is prohibitive. Table 3 investigates removing the vocabulary limit on the dictionary side, but keeping the continuous annotations vocabulary at 800 signs. We show that using the full 9k vocabulary from BSLDICT significantly improves the results on the unseen setting.

Batch Size. Next, we investigate the effect of increasing the number of negative pairs by increasing the batch size when training with Watch-Lookup on 1064 categories. We observe in Fig. 4(a) an improvement in performance with greater numbers of negatives before saturating. Our final Watch-Read-Lookup model has high memory requirements, for which we use 128 batch size. Note that the effective size of the batch with our sampling is larger due to sampling extra video clips corresponding to subtitles.

Temperature. Finally, we analyze the impact of the temperature hyperparameter τ on the performance of Watch-Lookup. We observe a major decrease in performance when τ approaches 1. We choose $\tau = 0.07$ used in [51, 55] for all other experiments. Additional ablations are provided in Appendix B.

BSL-1K Sign Spotting Benchmark. Although our learning framework primarily targets good performance on unseen continuous signs, it can also be naively applied to the (closed-vocabulary) sign spotting benchmark proposed by [8]. We evaluate the performance of our Watch-Read-Lookup model and achieve a score of 0.170 mAP, outperforming the previous state-of-the-art performance of 0.160 mAP [8].

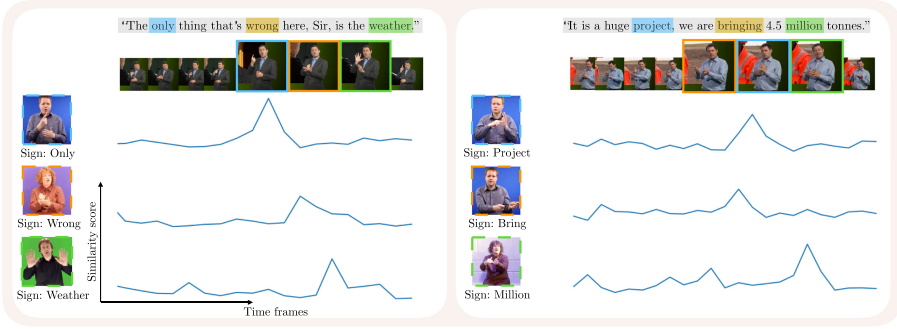


Fig. 6. Densification: We plot the similarity scores between BSL-1K test clips and BSLDICT examples over time, by querying only the words in the subtitle. The predicted locations of the signs correspond to the peak similarity scores.

4.4 Applications

In this section, we investigate three applications of our sign spotting method.

Sign Variant Identification. We show the ability of our model to spot specifically which variant of the sign was used. In Fig. 5, we observe high similarity scores when the variant of the sign matches in both BSL-1K and BSLDICT videos. Identifying such sign variations allows a better understanding of regional differences and can potentially help standardisation efforts of BSL.

Dense Annotations. We demonstrate the potential of our model to obtain dense annotations on continuous sign language video data. Sign spotting through the use of sign dictionaries is not limited to mouthings as in [8] and therefore is of great importance to scale up datasets for learning more robust sign language models. In Fig. 6, we show qualitative examples of localising multiple signs in a given sentence in BSL-1K, where we only query the words that occur in the subtitles, reducing the search space. In fact, if we assume the word to be known, we obtain 83.08% sign localisation accuracy on BSL-1K with our best model. This is defined as the number of times the maximum similarity occurs within $-20/+5$ frames of the end label time provided by [8].

“Faux Amis”. There are works investigating lexical similarities between sign languages manually [56,57]. We show qualitatively the potential of our model to discover similarities, as well as “faux-amis” between different sign languages, in particular between British (BSL) and American (ASL) Sign Languages. We retrieve nearest neighbors according to visual embedding similarities between BSLDICT which has a 9K vocabulary and WLASL [28], an ASL isolated sign language dataset, with a 2K vocabulary. We provide some examples in Fig. 7.



Fig. 7. “Faux amis” in BSL/ASL: Same/similar manual features for different English words (left), as well as for the same English words (right), are identified between BSLDICT and WLASL isolated sign language datasets.

5 Conclusions

We have presented an approach to spot signs in continuous sign language videos using visual sign dictionary videos, and have shown the benefits of leveraging multiple supervisory signals available in a realistic setting: (i) sparse annotations in continuous signing, (ii) accompanied with subtitles, and (iii) a few dictionary samples per word from a large vocabulary. We employ multiple-instance contrastive learning to incorporate these signals into a unified framework. Our analysis suggests the potential of sign spotting in several applications, which we think will help in scaling up the automatic annotation of sign language datasets.

Acknowledgements. This work was supported by EPSRC grant ExTol. The authors would like to thank A. Sophia Koepke, Andrew Brown, Necati Cihan Camgöz, and Bencie Woll for their help. S.A would like to acknowledge the generous support of S. Carlson in enabling his contribution, and his son David, who bravely waited until after the submission deadline to enter this world.

References

1. Sutton-Spence, R., Woll, B.: The Linguistics of British Sign Language: An Introduction. Cambridge University Press, London (1999)
2. Coucke, A., Chlieh, M., Gisselbrecht, T., Leroy, D., Poumeyrol, M., Lavril, T.: Efficient keyword spotting using dilated convolutions and gating. In: ICASSP (2019)
3. Véniat, T., Schwander, O., Denoyer, L.: Stochastic adaptive neural architecture search for keyword spotting. In: ICASSP (2019)
4. Momeni, L., Afouras, T., Stafylakis, T., Albanie, S., Zisserman, A.: Seeing wake words: audio-visual keyword spotting. In: BMVC (2020)
5. Stafylakis, T., Tzimiropoulos, G.: Zero-shot keyword spotting for visual speech recognition in-the-wild. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11208, pp. 536–552. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01225-0_32
6. Chung, J.S., Senior, A., Vinyals, O., Zisserman, A.: Lip reading sentences in the wild. In: CVPR (2017)
7. Afouras, T., Chung, J.S., Zisserman, A.: LRS3-TED: a large-scale dataset for visual speech recognition. arXiv preprint [arXiv:1809.00496](https://arxiv.org/abs/1809.00496) (2018)

8. Albanie, S., et al.: BSL-1K: scaling up co-articulated sign language recognition using mouthing cues. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020. LNCS, vol. 12356, pp. 35–53. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58621-8_3
9. Schembri, A., Fenlon, J., Rentelis, R., Cormier, K.: British Sign Language Corpus Project: A corpus of digital video data and annotations of British Sign Language 2008–2017 (Third Edition) (2017)
10. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: a new estimation principle for unnormalized statistical models. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 297–304 (2010)
11. Kadir, T., Bowden, R., Ong, E.J., Zisserman, A.: Minimal training, large lexicon, unconstrained sign language recognition. In: Proceedings of the BMVC (2004)
12. Tamura, S., Kawasaki, S.: Recognition of sign language motion images. *Pattern Recogn.* **21**, 343–353 (1988)
13. Starner, T.: Visual recognition of American sign language using hidden Markov models. Master’s thesis, Massachusetts Institute of Technology (1995)
14. Fillbrandt, H., Akyol, S., Kraiss, K.: Extraction of 3D hand shape and posture from image sequences for sign language recognition. In: IEEE International SOI Conference (2003)
15. Buehler, P., Everingham, M., Zisserman, A.: Learning sign language by watching TV (using weakly aligned subtitles). In: Proceedings of the CVPR (2009)
16. Cooper, H., Pugeault, N., Bowden, R.: Reading the signs: a video based sign dictionary. In: ICCVW (2011)
17. Ong, E., Cooper, H., Pugeault, N., Bowden, R.: Sign language recognition using sequential pattern trees. In: CVPR (2012)
18. Pfister, T., Charles, J., Zisserman, A.: Domain-adaptive discriminative one-shot learning of gestures. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8694, pp. 814–829. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10599-4_52
19. Agris, U., Zieren, J., Canzler, U., Bauer, B., Kraiss, K.F.: Recent developments in visual sign language recognition. *Univ. Access Inf. Soc.* **6**, 323–362 (2008)
20. Forster, J., Oberdörfer, C., Koller, O., Ney, H.: Modality combination techniques for continuous sign language recognition. In: *Pattern Recognition and Image Analysis* (2013)
21. Camgoz, N.C., Hadfield, S., Koller, O., Bowden, R.: SubUNets: end-to-end hand shape and continuous sign language recognition. In: ICCV (2017)
22. Huang, J., Zhou, W., Zhang, Q., Li, H., Li, W.: Video-based sign language recognition without temporal segmentation. In: AAAI (2018)
23. Ye, Y., Tian, Y., Huenerfauth, M., Liu, J.: Recognizing American sign language gestures from within continuous videos. In: CVPRW (2018)
24. Zhou, H., Zhou, W., Zhou, Y., Li, H.: Spatial-temporal multi-cue network for continuous sign language recognition. CoRR abs/2002.03187 (2020)
25. Camgoz, N.C., Koller, O., Hadfield, S., Bowden, R.: Sign language transformers: joint end-to-end sign language recognition and translation. In: CVPR (2020)
26. Carreira, J., Zisserman, A.: Quo Vadis, action recognition? A new model and the Kinetics dataset. In: CVPR (2017)
27. Joze, H.R.V., Koller, O.: MS-ASL: a large-scale data set and benchmark for understanding American sign language. In: BMVC (2019)
28. Li, D., Opazo, C.R., Yu, X., Li, H.: Word-level deep sign language recognition from video: a new large-scale dataset and methods comparison. In: WACV (2019)

29. Viitaniemi, V., Jantunen, T., Savolainen, L., Karppa, M., Laaksonen, J.: S-pot - a benchmark in spotting signs within continuous signing. In: LREC (2014)
30. Eng-Jon Ong, Koller, O., Pugeault, N., Bowden, R.: Sign spotting using hierarchical sequential patterns with temporal intervals. In: CVPR (2014)
31. Farhadi, A., Forsyth, D.A., White, R.: Transfer learning in sign language. In: CVPR (2007)
32. Bilge, Y.C., Ikizler, N., Cinbis, R.: Zero-shot sign language recognition: can textual data uncover sign languages? In: BMVC (2019)
33. Motiian, S., Jones, Q., Iranmanesh, S.M., Doretto, G.: Few-shot adversarial domain adaptation. In: NeurIPS (2017)
34. Zhang, J., Chen, Z., Huang, J., Lin, L., Zhang, D.: Few-shot structured domain adaptation for virtual-to-real scene parsing. In: ICCVW (2019)
35. Chang, W.G., You, T., Seo, S., Kwak, S., Han, B.: Domain-specific batch normalization for unsupervised domain adaptation. In: CVPR (2019)
36. Li, D., Yu, X., Xu, C., Petersson, L., Li, H.: Transferring cross-domain knowledge for video sign language recognition. In: CVPR (2020)
37. Koller, O., Forster, J., Ney, H.: Continuous sign language recognition: towards large vocabulary statistical recognition systems handling multiple signers. *Comput. Vis. Image Underst.* **141**, 108–125 (2015)
38. von Agris, U., Knorr, M., Kraiss, K.: The significance of facial features for automatic sign language recognition. In: 2008 8th IEEE International Conference on Automatic Face Gesture Recognition (2008)
39. Athitsos, V., et al.: The American sign language lexicon video dataset. In: CVPRW (2008)
40. Wilbur, R.B., Kak, A.C.: Purdue RVL-SLLL American sign language database. School of Electrical and Computer Engineering Technical report, TR-06-12, Purdue University, W. Lafayette, IN 47906 (2006)
41. Chai, X., Wang, H., Chen, X.: 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)
42. Schembri, A., Fenlon, J., Rentelis, R., Reynolds, S., Cormier, K.: Building the British sign language corpus. *Lang. Document. Conserv.* **7**, 136–154 (2013)
43. Cooper, H., Bowden, R.: Learning signs from subtitles: a weakly supervised approach to sign language recognition. In: CVPR (2009)
44. Chung, J.S., Zisserman, A.: Signs in time: encoding human motion as a temporal image. In: Workshop on Brave New Ideas for Motion Representations, ECCV (2016)
45. Dietterich, T.G., Lathrop, R.H., Lozano-Pérez, T.: Solving the multiple instance problem with axis-parallel rectangles. *Artif. Intell.* **89**, 31–71 (1997)
46. Pfister, T., Charles, J., Zisserman, A.: Large-scale learning of sign language by watching TV (using co-occurrences). In: BMVC (2013)
47. Feng, Y., Ma, L., Liu, W., Zhang, T., Luo, J.: Video re-localization. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. LNCS, vol. 11218, pp. 55–70. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01264-9_4
48. Yang, H., He, X., Porikli, F.: One-shot action localization by learning sequence matching network. In: CVPR (2018)
49. Cao, K., Ji, J., Cao, Z., Chang, C.Y., Niebles, J.C.: Few-shot video classification via temporal alignment. In: CVPR (2020)

50. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint [arXiv:1807.03748](https://arxiv.org/abs/1807.03748) (2018)
51. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018)
52. Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: End-to-end learning of visual representations from uncurated instructional videos. In: CVPR (2020)
53. <https://www.signbsl.com/>. (British sign language dictionary)
54. Cao, Z., Hidalgo, G., Simon, T., Wei, S.E., Sheikh, Y.: OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In: arXiv preprint [arXiv:1812.08008](https://arxiv.org/abs/1812.08008) (2018)
55. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)
56. SignumMcKee, D., Kennedy, G.: Lexical comparison of signs from American, Australian, British and New Zealand sign languages. An anthology to honor Ursula Bellugi and Edward Klima, *The signs of language revisited* (2000)
57. Aldersson, R., McEntee-Atalians, L.: A lexical comparison of Icelandic sign language and Danish sign language. *Birkbeck Stud. Appl. Ling.* **2**, 123–158 (2007)