

Weakly-supervised Fingerspelling Recognition in British Sign Language Videos

K R Prajwal*¹

<http://robots.ox.ac.uk/~prajwal>

Hannah Bull*²

<https://hannahbull.github.io>

Liliane Momeni*¹

<https://www.robots.ox.ac.uk/~liliane>

Samuel Albanie³

<https://samuelalbanie.com>

Gül Varol⁴

<http://imagine.enpc.fr/~varolg>

Andrew Zisserman¹

<http://robots.ox.ac.uk/~az>

¹ Visual Geometry Group

Department of Engineering Science

University of Oxford

Oxford, UK

² LISN, Univ Paris-Saclay,

CNRS, France

³ Department of Engineering, University of
Cambridge, UK

⁴ LIGM, École des Ponts, Univ Gustave
Eiffel, CNRS, France

Abstract

The goal of this work is to detect and recognize sequences of letters signed using fingerspelling in British Sign Language (BSL). Previous fingerspelling recognition methods have not focused on BSL, which has a very different signing alphabet (e.g., two-handed instead of one-handed) to American Sign Language (ASL). They also use manual annotations for training. In contrast to previous methods, our method only uses weak annotations from subtitles for training. We localize potential instances of fingerspelling using a simple feature similarity method, then automatically annotate these instances by querying subtitle words and searching for corresponding mouthing cues from the signer. We propose a Transformer architecture adapted to this task, with a multiple-hypothesis CTC loss function to learn from alternative annotation possibilities. We employ a multi-stage training approach, where we make use of an initial version of our trained model to extend and enhance our training data before re-training again to achieve better performance. Through extensive evaluations, we verify our method for automatic annotation and our model architecture. Moreover, we provide a human expert annotated test set of 5K video clips for evaluating BSL fingerspelling recognition methods to support sign language research.

1 Introduction

Fingerspelling in signed languages is a means to encode words from the surrounding written language into sign language via a manual alphabet, i.e. one sign per letter. Words from a written language with no known sign may be fingerspelled, such as names of people and places. Additionally, some signs are derived from fingerspelled words, for example, initialized signs in BSL such as ‘F’ for ‘father’ or ‘KK’ for ‘kitchen’ [19]. Padden & Gunsauls [25] estimate that signers fingerspell 12-35% of the time in ASL. Within the BSL videos used in this work [8], we estimate roughly 5-10% based on the duration of automatically detected fingerspelling content. Consequently, it




Figure 1: Fingerspelling recognition. We study the task of recognising a sequence of BSL finger-spelled letters in a continuous signing window in sign language interpreted TV broadcast data. We exploit the accompanying English subtitles to automatically collect training data. However, the task remains very challenging as (i) the signer may fingerspell only a subset of the letters – as shown above, although the subtitle contains the word ‘Astley’, only the letters ‘A’, ‘S’ and ‘Y’ are finger-spelled – and (ii) due to occlusions from the hands, different letters can visually look very similar.

is important to incorporate automatic fingerspelling recognition methods to be able to *exhaustively* transcribe signs in continuous sign language videos.

Some sign languages, such as ASL and LSF, use a one-handed manual alphabet, and others, including BSL and Auslan, use a two-handed manual alphabet [40]. One-handed manual alphabets typically do not involve significant wrist movements, while two-handed fingerspelling resembles other lexical signs, making it relatively more difficult to detect fingerspelling beginning and end times from a longer signing sequence. The presence of two-hand movements further results in occlusions, making it challenging to differentiate between certain characters, e.g. especially vowels.

In this work, we focus on BSL, i.e. the more challenging two-handed fingerspelling. Given a sign language video, our goal is to temporally locate the fingerspelling segments within the video (i.e. detection) and to transcribe the fingerspelled letters to text (i.e. recognition). To this end, we design a Transformer-based model that jointly performs both detection and recognition. Our key contribution lies in the data collection procedure used to automatically obtain training data for this task, which is applicable to any sign language videos that have approximately-aligned subtitle translations. We also provide the first large-scale benchmark for BSL fingerspelling recognition based on the recently released BOBSL dataset [9]. Our experiments on this benchmark demonstrate promising results with a 53.3 character error rate on this challenging task while only using weak supervision.

Previous works building fingerspelling datasets rely on manual annotation, either by expert annotators resulting in limited data [42], or by crowdsourcing noisy large-scale annotations [43]. In contrast to these works on ASL fingerspelling, we introduce a practical methodology to automatically annotate fingerspelling in the presence of subtitled sign language video data, allowing to scale up the data size, and potentially to be applicable to other sign languages. Starting from a small number of manually annotated fingerspelling exemplars, we use an embedding space to find numerous similar instances of fingerspelling in the corpus. To annotate these instances, we exploit the observation that signers often simultaneously mouth the words (i.e. silent speech with lip movements) which they fingerspell. We obtain an initial set of annotations by querying potential words from the subtitles, especially proper nouns, and identifying mouthing cues [4, 44] which also coincide with fingerspelling instances. This initial set of annotations are further extended and enhanced by a pseudolabeling step, also making use of subtitles (see Sec. 4).

A key challenge with training through automatic annotations is label noise. For example, the mouthing model [44] may spot a single word, while the fingerspelling contains multiple words, such as name-surname pairs. In some other cases, the mouthing model may fail, associating the wrong word in the subtitle with the fingerspelling segment. Also, some letters of the word may be skipped within fingerspelling as illustrated in Fig. 1. To account for this uncertainty, we implement a multiple-hypotheses version of the CTC loss [45] (MH-CTC) where we consider all nouns from

the subtitle, as well as bigrams and trigrams, as potential targets.

Our main contributions are: (i) Training a BSL fingerspelling detection and recognition model using only weak labels from subtitles, mouthing cues, and a small number (115) of manual fingerspelling exemplars; (ii) Employing multiple hypotheses from the subtitle words within the CTC loss to train with noisy labels; (iii) Demonstrating advantages of a pseudolabeling step incorporating the subtitle information; and (iv) Providing a large-scale manually annotated benchmark for evaluating BSL fingerspelling recognition, released for research purposes. Please check our website for more details: <https://www.robots.ox.ac.uk/~vgg/research/transpeller>.

2 Related work

Our work relates to four themes in the research literature: the broader topics of *sign language recognition* and *learning sign language from weak/noisy annotation*, and the more directly related literature on *spotting mouthings* and *fingerspelling recognition*.

Sign language recognition. Building on the pioneering 1988 work of Tamura and Kawasaki [60], early approaches to automatic sign recognition made use of hand-crafted features for motion [65] and hand shape [18, 64]. To model the temporal nature of signing, there has also been a rich body of work exploring the use of Hidden Markov Models [10, 17, 25, 26, 47, 63] and Transformers [2, 12]. One notable trend in prior work is the transition towards employing deep spatiotemporal neural networks to provide robust features for recognition. In this regard, the I3D model of Carreira and Zisserman [8] has seen widespread adoption, achieving strong recognition results on a range of benchmarks [2, 22, 27, 28]. In this work, we likewise build our approach on strong spatiotemporal video representations, adopting the Video-Swin Transformer [60] as a backbone for our model.

Learning sign language from weak/noisy annotation. Given the paucity of large-scale annotated sign language datasets, a range of prior work has sought to leverage weakly aligned subtitled interpreter footage as a supervisory signal [6, 9, 33, 56] for sign spotting and recognition via apriori mining [10] and multiple instance learning [14]. Similarly to these works, we likewise aim to make use of subtitled signing footage. However, we do so in order to detect and recognize fingerspelling in a manner that allows for scalable training. To the best of our knowledge, this approach has not been considered in prior work.

Spotting mouthings in sign language videos. Signers often mouth the words that they sign (or fingerspell) [18]. The recent advancements in visual keyword spotting [82, 67, 46] have enabled the automatic curation of large-scale sign language datasets by spotting a set of query words using the mouthing cues and matching them with the corresponding sign segment. The state-of-the-art architecture for the visual KWS task is the Transpotter [67], which we use in Sec 4.1 to obtain our initial set of automatic fingerspelling annotations. Our fingerspelling architecture also partly takes inspiration from the Transpotter (and more broadly from prior works for text spotting that detect words and learn to read them via CTC [2, 29]), wherein we process video features with a single Transformer encoder and then employ multiple heads to solve related tasks such as detection and classification with a [CLS] token. In this work, we add also add a recognition head supervised by a novel loss function. We also show the benefits of our multi-stage training pipeline in training this model when we only have weak supervision.

Fingerspelling recognition and detection. Early work on automatic fingerspelling recognition explored the task of classification under fairly constrained settings, focusing on isolated signs and limited vocabularies (e.g. 20 words [20], 82 words [69] and 100 words [61]). Kim et al. propose to consider instead a “lexicon-free” setting which they tackle with frame-level classifiers in combination with segmental CRFs on a newly introduced dataset, of 3,684 American Sign Language (ASL) fingerspelling instances [22]. Moving towards more challenging data, the ChicagoFSWild (7304 fingerspelling sequences across 160 signers annotated by ASL students) [42] and ChicagoFSWild+ (55,232 fingerspelling sequences signed by 260 signers annotated by crowdsourcing) [43] datasets




Figure 2: **Transpeller architecture.** Given a short clip of a signer, we extract features from a visual front-end pre-trained on the sign language recognition task. We project these features to a desired hidden dimension and add positional encodings before encoding these feature vectors using a transformer encoder. We use three heads on top of the transformer outputs to predict if the clip contains fingerspelling (classification prediction), and if so, where it is in the clip (localization prediction) and what it is (recognition prediction). The subtitle shown on the left is used for curating the training data and is shown here for illustrative purposes.

sourced from YouTube and Deaf social media target greater diversity and visual variation.

From a modeling perspective, Pugealt and Bowden employ random forests on depth and intensity images for real-time recognition of 24 fingerspelled letters [38]. Shi et al. demonstrate the benefits of using a signing hand detector for fingerspelling recognition without frame-level labels [40], motivating later work to attain this benefit automatically through visual attention without an explicit region detector [43]. Other work has explored the feasibility of using synthetic hand training data to fine-tune a CNN for isolated Irish Sign Language (ISL) fingerspelling recognition [49].

More closer to our work, several works have considered detecting the temporal location of fingerspelling in addition to recognition. This includes efforts to segment signing into sign types (classifiers, lexical signs, and fingerspelling) prior to recognition [56], as well as systems for fingerspelling detection supervised with segment boundaries [44]. The recently proposed FSS-Net learns joint embeddings to enable fingerspelling search within and across videos [45]. In contrast, our work is weakly supervised with noisy, automatic annotations. It is weakly supervised in the sense that the model lacks access to ground truth fingerspelling boundaries at training, while the annotation is noisy in the sense that it is derived from subtitles from which the signing is produced as a translation, rather than a transcription.

3 Fingerspelling detection and recognition

In this section, we introduce Transpeller, our Transformer-based model to recognize and detect fingerspelling (Sec. 3.1). Our model is trained only on automatically curated data. In order to circumvent this label noise, we also propose a new loss function in Sec. 3.2.

3.1 The Transpeller architecture

Our model ingests a video clip of a signer, encodes it with a Transformer encoder [57], and produces three outputs with each of its prediction heads: (i) a classification head that predicts **if** the given video clip contains fingerspelling, (ii) a localization head that produces per-frame probabilities indicating **where** the fingerspelling is in the clip, (iii) a recognition head that produces a sequence of letter probabilities indicating **what** is being fingerspelled. We illustrate this architecture in Figure 2.

Visual backbone. Our input is a sequence of RGB frames constituting a short clip of a signer. In order to extract the visual features, we follow prior sign-language works [6, 64, 61] and use a pre-trained sign classification model. The sign classification model used in prior works is an I3D model [6, 64, 61] pre-trained on Kinetics [23] and finetuned using a sign classification objective. We replace the I3D with a relatively modern Video-Swin-S [60] architecture and finetune it for

the sign classification task in a similar fashion. We pre-extract the features for all our videos and save them at a temporal stride of 4. The Transpeller operates on these feature vectors $V \in \mathbb{R}^{T \times d_f}$. **Transformer Encoder.** Given a sequence of visual feature vectors $V \in \mathbb{R}^{T \times d_f}$, we first use two fully-connected layers (FC) to project the feature dimension to d , which is the hidden dimension of the transformer encoder. We add the temporal positional encodings and prepend this sequence with a learnable $[CLS]$ token embedding (such as in BERT [13] and ViT [14]). We encode the temporal information of this sequence by passing it through a Transformer encoder consisting of N layers:

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

Prediction heads. The $[CLS]$ output feature vector $V_{enc(1)}$ serves as a aggregate representation for the entire input video clip. An MLP head for binary classification, f_c is attached to $V_{enc(1)}$ to predict the probability of a fingerspelling segment being present in the input video:

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

where σ denotes a sigmoid activation. To localize the fingerspelling segment, we attach a second MLP head f_l that is shared across the encoded video feature time-steps $V_{enc(2:T+1)}$:

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

The output y_t^{loc} at each feature time-step $t \in T$ indicates the probability of the time-step t being a part of a fingerspelling segment. In order to recognize what is being fingerspelt, a third MLP head f_r is attached in a similar fashion to f_l to predict \mathbb{C} letter probabilities at each time-step:

$$\hat{y}^{rec} = \text{softmax}(f_r(V_{enc(2:T+1)})) \in \mathbb{R}^{T \times \mathbb{C}}.$$

Loss functions. Given a training dataset \mathcal{D} consisting of video clips v , class labels y^{cls} , location labels y^{loc} and recognition labels (whenever present) y^{rec} , we define the following training objectives:

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

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

$$\mathcal{L}^{rec} = -\mathbb{E}_{(v, y^{rec}) \in \mathcal{D}} \text{CTC}(y^{rec}, \hat{y}^{rec}) \quad (3)$$

where BCE stands for the binary cross-entropy loss and CTC stands for Connectionist Temporal Classification [24] loss. The labels y^{cls} are set to 1 when the given keyword occurs in the video and 0 otherwise; the frame labels y^{loc} are set to 1 for the frames where the keyword is uttered and 0 otherwise. We optimize the total loss $\mathcal{L} = \mathcal{L}^{cls} + \mathcal{L}^{loc} + \lambda \mathcal{L}^{rec}$ where, $\lambda = 1$ if a given input video has a word label annotation, else, $\lambda = 0$ to only train the classification and localization heads.

3.2 Multiple Hypotheses (MH) CTC loss

As described earlier, our word labels for fingerspelling recognition are weak labels from subtitles and mouthing cues. The process for obtaining these labels is described in Sec. 4.1 & 4.2. Since this process is automatic, it introduces a degree of label noise. For example, the mouthing model can detect false positives; the detection boundaries can be erroneous, or multiple mouthings can be spotted for a single detection interval. We observed that our automatic detection pipeline is more accurate than our process of obtaining word labels. If we assume the fingerspelling detection is correct, then it is quite likely that the fingerspelled word is among one of the words (usually a noun) in the subtitle.

With this idea in mind, we design an improved CTC loss function, termed Multiple Hypotheses CTC (MH-CTC), that allows the model to “pick” the most correct word label from a set of possible word hypotheses. This set comprises a number of words, of which one is the correct word label for the fingerspelling sequence in the input video clip. For example, the set of word hypotheses

	Train			Val		
	#labels	vocab	avg. dur.	#labels	vocab	avg. dur.
Stage 1: Exemplar detections	149k	-	1.4s	3.0k	-	1.3s
w/ mouthing labels: nouns	59k	18k	1.8s	1.2k	0.8k	1.6s
w/ mouthing labels: pr. nouns	39k	14k	1.9s	0.9k	0.5k	1.7s
Stage 2: Transpeller detections	129k	-	1.9s	2.5k	-	1.8s
w/ mouthing labels: nouns	61k	19k	2.1s	1.3k	0.9k	1.9s
w/ mouthing labels: pr. nouns	41k	15k	2.2s	0.9k	0.5k	2.0s
w/ Transpeller labels	111k	32k	2.0s	2.2k	1.4k	1.8s

Table 1: **Two stages of automatic annotations.** Stage 1: We use exemplars to detect fingerspelling and mouthing cues to obtain letter labels. Stage 2: We obtain detections and letter labels using Transpeller pseudolabels.

could be the proper nouns in the subtitle corresponding to the input video clip. Given a list of word hypotheses \mathbb{H} for the input video clip v , our modified recognition loss \mathcal{L}^{rec} is:

$$\mathcal{L}^{rec} = -\mathbb{E}_{(v, \mathbb{H}) \in \mathcal{D}} \min_{\forall h \in \mathbb{H}} CTC(h, \hat{y}^{rec})$$

This corresponds to backpropagating the recognition loss for the word that achieves the minimum CTC loss among all the hypotheses. Since this allows the model to “choose” its own target, we found two strategies that help the model converge. Firstly, pretraining the model with CTC loss using Eqn 3 before using MH-CTC is essential. Secondly, when using MH-CTC, we found that randomly (with 50% chance) setting \mathbb{H} to be a single hypothesis containing the word found by our automatic annotation also prevents the model from diverging.

4 Automatic annotations

Our model is completely trained with weak labels that are automatically curated. We perform two stages of automatic annotation. In the first stage (Sec. 4.1), we automatically annotate fingerspelling detections using exemplars and letter labels using mouthing cues. In the second stage (Sec. 4.2), we obtain detection and letter pseudolabels from the Transpeller model that has been trained on annotations from the first stage. The number and duration of these fingerspelling detections, as well as the number of detections associated with either a noun or proper noun letter label, is shown in Tab. 1.

4.1 Exemplar and mouthing annotations

Detections using exemplars. We manually annotate a small number E of single frame exemplars of fingerspelling amongst videos from different signers in the training set ($E = 115$). Using these exemplars, we search for frames in continuous signing with high feature similarity to these frames containing fingerspelling. This exemplar-based annotation technique is inspired by [63, 64]. We use features from [64] and compute cosine similarity with fingerspelling features. This simple method provides approximate annotation of fingerspelling detections. We use this method for two reasons: firstly, to help our model learn fingerspelling temporal detection, and secondly, to select video segments containing fingerspelling for manual annotation. Technical details on the computation of the feature similarity can be found in the supplementary material.

Letter labels from mouthings. To obtain word annotations for the fingerspelled segments, we use mouthing cues, as fingerspelled words are often mouthed simultaneously in interpreted data [41]. In [64], the authors use an improved Transpotter architecture [62] to query words from surrounding subtitles and localize corresponding mouthing cues. We consider all mouthing annotations from [64] falling within the interval of the fingerspelling detections. As the fingerspelling detection boundaries are approximate, and the automatic mouthing annotations are not always accurate, these annotations

are noisy. Tab. 2 shows that almost all fingerspelling annotations refer to nouns, and most refer to proper nouns. Thus, restricting mouthing annotations to nouns or proper nouns reduces noise.

4.2 Transpeller annotations

Improving detections with Transpeller pseudolabels. The model described in Sec. 3.1 outputs a per-frame localization score, predicting the presence of fingerspelling. After training this model, we can improve upon the exemplar-based detections using pseudolabels. Given localisation scores s_1, \dots, s_N for a window of N frames, we consider that the window contains fingerspelling if $\max(s_1, \dots, s_N) > t_1$. We consider that a sub-interval $[i, j]$ ($1 \leq i \leq j \leq N$) of this window contains fingerspelling if $s_i, \dots, s_j \geq t_2$, where $t_2 < t_1$. To smooth the localization scores, we take the moving maximum score amongst K consecutive scores. We let $t_1 = 0.7$, $t_2 = 0.3$ and $K = 5$.

Improving letter labels with Transpeller pseudolabels. Using pseudolabels from the model in Sec. 3.1, we can also improve automatic letter label annotations. After decoding the CTC outputs using beam search, we can then compute a proximity score with words from neighboring subtitles, i.e. $\text{dist}(w_1, w_2)$, where w_1 is a subtitle word and w_2 is the output of beam search decoding. The proximity score is a variant of the Levenshtein edit distance, but where deletions are not heavily penalized. This is because words such as ‘Sarah Jane’ can be reasonably fingerspelled as ‘SJ’. Details of this proximity score are provided in the supplementary material. We can use this proximity score to find the subtitle word most likely to be fingerspelled.

4.3 Multi-stage training

We perform a multi-stage training strategy where we start by training on exemplar-based annotations (lines 2 and 3 of Tab. 1) using the vanilla CTC loss to supervise the recognition head. Upon convergence, we finetune this model further using MH-CTC, this time additionally considering proper nouns in neighboring subtitles as the hypotheses.

Using the above pre-trained model, we now extract the Transpeller annotations as detailed in Sec. 4.2. Using these new annotations, we repeat the process: we start with the vanilla CTC loss and then finetune this model further using MH-CTC. Given that we have pseudo-labels from the Stage 1 Transpeller model, we can restrict the MH-CTC search space. The hypotheses now only contain nouns and proper nouns from neighboring subtitles with at least one letter in common with the CTC decoded outputs from the Transpeller model of Stage 1.

5 Experiments

5.1 BOBSL Fingerspelling benchmark

We collect and release the first benchmark for evaluating fingerspelling in British Sign Language. The test set annotations are collected by adapting the VIA Whole-Sign Verification Tool [16] to the task of fingerspelling recognition. Given proposed temporal windows around the automatic exemplar and mouthing annotations, annotators mark whether there is any fingerspelling in the signing window and type out the exact letters which are fingerspelled. We use a temporal window of 2.1s before to 4s after the midpoint of the automatically detected fingerspelling instance. Since the fingerspelled letters could be a subset of the actual full word (e.g. SH for SARAH), we also obtain the corresponding full word annotations. Descriptive statistics on the test set annotations are in Tab. 2.

Evaluation criteria. We measure the fingerspelling recognition performance using the Character Error Rate (CER), which provides a normalized count of the substitution/deletion/insertion errors in the predicted letter sequence when compared to the ground-truth sequence. We report two CERs for the two different ground-truth annotations we have for each clip: (i) $\text{CER}_{\text{fspell}}$ - ground-truth is the actual fingerspelled letters which, as mentioned before, may only be part of a word, and (ii) CER_{full} - ground-truth is the full word annotation to which the fingerspelling refers to. Given that our model(s) are trained only with automatic weakly-supervised full-word annotations (Sec. 4),

#labels	(full word) vocab.	%nouns	%pr. nouns	%full	avg. % missing
4923	3442	96%	74%	22%	34%

Table 2: **Statistics on test set annotations.** Most fingerspellings refer to proper nouns or nouns. Around 22% of fingerspelling instances contain all letters of the encoded word, but on average 34% of the letters of a word are omitted during fingerspelling.

Annotations	# Recognition ex.	CER _{fspell}	CER _{full}
Exemplars + Mouthings: proper nouns	39k	58.5	62.1
Exemplars + Mouthings: nouns	59k	58.6	62.9
finetuned with MH-CTC	59k (avg. 4 hypoth.)	57.6	64.3

Table 3: **Stage 1: Transpeller model with exemplar + mouthing supervision.** All rows use 149k fingerspelling detections for training. When assigning a word label for these detections from the subtitle, we choose to look at nouns, especially proper nouns, as they are more likely to be fingerspelt. We obtain the best results when using MH-CTC loss.

Annotations	# Recognition ex.	CER _{fspell}	CER _{full}
Transpeller detect. + Mouthings	61k	57.5	63.1
Transpeller detect. + Char. labels	111k	55.4	63.0
finetuned with MH-CTC	111k + (avg. 9 hypoth.)	53.3	60.1

Table 4: **Stage 2: Transpeller pseudolabels.** All rows use 129k fingerspelling detections for training. Using both the refined detections and word annotations from the Stage 1 model gives a clear reduction in the CER. Further, using MH-CTC gives a 2.1 CER boost. Overall, our Stage 2 achieves a final best CER of 53.3 which is 4.3 CER better than the best model of Stage 1, thus validating the impact of our multi-stage training pipeline.

these two different CER scores can help us see if the model learns to pick up on the fingerspelled letters, rather than only relying on the mouthing cues or memorizing the full word annotations.

Implementation details. We use a batch size of 32 and an initial learning rate of $5e^{-5}$, which is reduced to $1e^{-5}$ after the validation loss does not improve for 3 epochs. At test time, we decode with a beam width of 30. More implementation details are provided in the supplementary material.

5.2 Results

We now evaluate different variations at each stage of the Transpeller recognition pipeline.

Using exemplar detections and mouthing cue letter labels with CTC loss. Our initial set of annotations from the exemplar detections gives us 149k fingerspelling instances, out of which a fraction of them can be associated with word labels using mouthing cues. In Tab. 3, we show how our performance depends on our choice of recognition annotations. Restricting our automatic word label annotations to proper nouns gives us a cleaner training set, as they are most likely to be fingerspelt. However, this comes at the expense of having very few training samples. In row 2 of Table 3, we find that we can tolerate a bit of label noise and expand to all nouns. We finetune the best CTC-based model from the above using our MH-CTC loss, which further results in an improvement of 1.0 CER_{fspell}. This is our best model with the initial set of annotations.

Transpeller detections. We now expand our training data, by extracting pseudolabels (Sec. 4.2) from the best model from the previous stage. We train on these new annotations and report our results in Tab 4. When we use the refined detections but still use the mouthing cues for assigning word

labels, we obtain a $\text{CER}_{\text{fspell}}$ of 57.5, a similar result to the corresponding model (57.6) from Stage 1. **Transpeller letter labels.** The error rates drop further when we also improve the recognition annotations. We do so by using a variant of the edit distance to match the predictions of the Stage 1 Transpeller to a word in a neighboring subtitle. These results can be seen in row 2 of Tab. 4.

Finetuning with MH-CTC. Finally, when finetuning further with our MH-CTC loss, our final Stage 2 model gives a $\text{CER}_{\text{fspell}}$ of 53.3 which is 4.3 points better than the best model (57.6) of Stage 1. It is evident from both these tables that both MH-CTC and our multi-stage training with pseudolabelling improve our recognition performance.

Fingerspelling recognition vs. spotting mouthings. An interesting line of thought is: if we have access to the subtitles at test time, can we use a mouthing model to accurately predict words instead of doing fingerspelling recognition? To judge this, we restrict the test set to instances where there is a noun or a proper noun mouthing annotation and compute the CER between the mouthing annotation and the ground truth annotation, and we get 55.6. On this same subset, our best model obtains 53.3 CER, demonstrating that it in fact performs better than a mouthing method with access to the subtitle text.

Error analysis of the best model on the test set. In Tab. 5, we show a few examples of our predictions and the corresponding ground-truth sequence. We can see our model makes reasonable errors for most examples, where it confuses between letters that are visually quite similar. In Fig. 4, we show how the CER varies based on the length of the ground-truth character sequence. It is evident that the model struggles the most with very short sequences (one or two letters). This is expected because it has only been supervised with full word annotations during training and has never been trained to predict one or two letters in isolation.

Lookup-based correction at inference-time. We explore the possibility of correcting the errors in the model’s outputs at test-time with the help of a pre-determined list of words, e.g. “atlantic” to “atlantic”. We first curate a list of nouns present in the subtitles of the BOBSL train set and use edit distance to match the model’s predictions to the closest noun in our list. If the edit distance is below a set threshold, i.e., a very close match, we replace the predicted character sequence with the matched word from our list. However, we found that this increases the $\text{CER}_{\text{fspell}}$ to 57.0 and CER_{full} to 61.1. Such a lookup-based correction method leads to several false matches because (i) it is done with no context of the surrounding words, (ii) not all the letters of the word are fingerspelled, (iii) fingerspelled words in the test set can also be novel and unseen.

5.3 Architecture ablations

Importance of joint recognition and detection. As described in Sec 3.1, our model contains three prediction heads for classification, localization, and recognition. All three heads are essential and are used to obtain the Stage 2 annotations as described in 4.2. We conduct an experiment to also demonstrate that it is beneficial to train these heads jointly. We find that joint training leads to a better recognition performance (55.4 $\text{CER}_{\text{fspell}}$) than training without the localization head (56.3 $\text{CER}_{\text{fspell}}$) or without the localization and classification heads (56.2 $\text{CER}_{\text{fspell}}$).

Sequence-to-Sequence vs. CTC-based models. We also compare our CTC-based recognition head with a sequence-to-sequence (seq2seq) encoder-decoder architecture supervised with a cross-entropy loss. We use the standard Transformer-Base [52] model, which contains a Transformer encoder similar to Transpeller and a 6-layer auto-regressive Transformer decoder. We compare with the CTC model trained on the Stage 2 annotations. The seq2seq network performs worse (57.4 $\text{CER}_{\text{fspell}}$) than the CTC model (55.4 $\text{CER}_{\text{fspell}}$). This is expected because only 22% of the test samples contain no missing letters, i.e. full words, so conditioning on past letters can lead to errors in future letter predictions. However, conditioning on past letters can be very helpful if we want to actually estimate the full word and not determine the exact letters that are fingerspelled. This is indeed validated by the fact that the seq2seq model performs much better on CER_{full} , achieving 55.2 compared to the CTC model’s 63.0. We also note that due to the sequential decoding in the




Figure 3: **Benefit of Transpeller annotations:** Here, we are given a fingerspelling clip of “BERKELEY”. The initial annotation with exemplars + mouthing cues assigns an incorrect word label because the spotted mouthing of the word “Homes” is temporally close to the fingerspelling location, resulting in a false word assignment. In the second annotation stage, we correct this using the Transpeller’s predicted characters “BEREY” that are matched to the subtitle word “BERKELEY”.




Figure 4: **Variation of CER vs the number of letters in the ground-truth.** Transpeller struggles to correctly predict very short fingerspelling segments, which are usually partial fingerspellings (e.g. MJ for Mary Jane). This happens because the model is only supervised with full-length words during training.

Ground Truth	Prediction	CER
chloride	churide	25
nuremburg	turmug	50
ivory	ener	80
elind	elinc	20
bnmm	ben	75
clove	clune	40
semnoa	samol	50

Table 5: **Qualitative examples.** Model predictions on our manually verified test set. The CERs are shown for reference. We see that one of the most common error sources is confusion between letters that are visually similar: (a, e, i), (d, c), (l, n, m, v, t), (o, u).

seq2seq model, it is $\approx 20\times$ slower in terms of run-time speed than the CTC model, which decodes all the characters in parallel.

6 Conclusion

We presented a new BSL fingerspelling recognition benchmark and our Transpeller model designed to jointly detect and recognize fingerspelled letters in continuous sign language video. Our training data is largely constructed automatically, exploiting English subtitles and mouthing cues. The evaluation data is manually curated, and we achieve promising recognition results (Tab. 5). However, we note some limitations. First, there remains room for improvement in the accuracy of the model to further reduce the character error rate. Second, our training and evaluation of the Transpeller is limited to the use of interpreted data and therefore is not necessarily representative of more natural, conversational signing. Addressing these limitations would be a valuable future research direction.

Acknowledgements

This work was supported by the Oxford-Google DeepMind Graduate Scholarship, EPSRC grant Ex-Tol, a Royal Society Research Professorship and the ANR project CorVis ANR-21-CE23-0003-01. HB would like to thank Annelies Braffort and Michèle Gouiffès for the support.

References

- [1] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. Mining association rules between sets of items in large databases. In *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, pages 207–216, 1993.
- [2] Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Joon Son Chung, Neil Fox, and Andrew Zisserman. Bsl-1k: Scaling up co-articulated sign language recognition using mouthing cues. In *European conference on computer vision*, pages 35–53. Springer, 2020.
- [3] Samuel Albanie, Gül Varol, Liliane Momeni, Hannah Bull, Triantafyllos Afouras, Himel Chowdhury, Neil Fox, Bencie Woll, Rob Cooper, Andrew McParland, and Andrew Zisserman. BOBSL: BBC-Oxford British Sign Language dataset. *arXiv*, 2021.
- [4] Fedor Borisjuk, Albert Gordo, and Viswanath Sivakumar. Rosetta: Large scale system for text detection and recognition in images. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 71–79, 2018.
- [5] Patrick Buehler, Andrew Zisserman, and Mark 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.
- [6] Hannah Bull, Triantafyllos Afouras, Gül Varol, Samuel Albanie, Liliane Momeni, and Andrew Zisserman. Aligning subtitles in sign language videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11552–11561, 2021.
- [7] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. Sign language transformers: Joint end-to-end sign language recognition and translation. In *CVPR*, 2020.
- [8] Joao Carreira and Andrew 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.
- [9] Helen Cooper and Richard Bowden. Learning signs from subtitles: A weakly supervised approach to sign language recognition. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2568–2574. IEEE, 2009.
- [10] Helen Cooper, Nicolas Pugeault, and Richard Bowden. Reading the signs: A video based sign dictionary. In *2011 IEEE international conference on computer vision workshops (ICCV Workshops)*, pages 914–919. IEEE, 2011.
- [11] Jeffrey Davis. Linguistic transference and interference: Interpreting between. *Sign language research: Theoretical issues*, page 308, 1990.
- [12] Mathieu De Coster, Mieke Van Herreweghe, and Joni Dambre. Sign language recognition with transformer networks. In *12th international conference on language resources and evaluation*, pages 6018–6024. European Language Resources Association (ELRA), 2020.

- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [14] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [16] Abhishek Dutta and Andrew Zisserman. The via annotation software for images, audio and video. In *Proceedings of the 27th ACM international conference on multimedia*, pages 2276–2279, 2019.
- [17] Gaolin Fang, Xiujuan Gao, Wen Gao, and Yiqiang Chen. A novel approach to automatically extracting basic units from chinese sign language. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 4, pages 454–457. IEEE, 2004.
- [18] Holger Fillbrandt, Suat 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.
- [19] Frank Fowley and Anthony Ventresque. Sign language fingerspelling recognition using synthetic data. In *AICS*, pages 84–95, 2021.
- [20] Paul Goh and Eun-Jung Holden. Dynamic fingerspelling recognition using geometric and motion features. In *2006 International Conference on Image Processing*, pages 2741–2744. IEEE, 2006.
- [21] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006.
- [22] Hamid Reza Vaezi Joze and Oscar Koller. MS-ASL: A large-scale data set and benchmark for understanding american sign language. In *BMVC*, 2019.
- [23] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [24] Taehwan Kim, Jonathan Keane, Weiran Wang, Hao Tang, Jason Riggle, Gregory Shakhnarovich, Diane Brentari, and Karen Livescu. Lexicon-free fingerspelling recognition from video: Data, models, and signer adaptation. *Computer Speech & Language*, 46: 209–232, 2017.
- [25] Oscar Koller, O Zargaran, Hermann Ney, and Richard Bowden. Deep sign: Hybrid cnn-hmm for continuous sign language recognition. In *Proceedings of the British Machine Vision Conference 2016*, 2016.
- [26] Oscar Koller, Sepehr Zargaran, and Hermann Ney. Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent cnn-hmms. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4297–4305, 2017.

- [27] Dongxu Li, Cristian Rodriguez Opazo, Xin Yu, and Hongdong Li. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *WACV*, 2019.
- [28] Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, and Hongdong Li. Transferring cross-domain knowledge for video sign language recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6205–6214, 2020.
- [29] Hui Li, Peng Wang, and Chunhua Shen. Towards end-to-end text spotting with convolutional recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5238–5246, 2017.
- [30] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3202–3211, 2022.
- [31] Stephan Liwicki and Mark Everingham. Automatic recognition of fingerspelled words in british sign language. In *2009 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 50–57. IEEE, 2009.
- [32] Liliane Momeni, Triantafyllos Afouras, Themis Stafylakis, Samuel Albanie, and Andrew Zisserman. Seeing wake words: Audio-visual keyword spotting. *arXiv preprint arXiv:2009.01225*, 2020.
- [33] Liliane Momeni, Gül Varol, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Watch, read and lookup: Learning to spot signs from multiple supervisors. In *Asian Conference on Computer Vision*, 2020.
- [34] Liliane Momeni, Hannah Bull, K R Prajwal, Samuel Albanie, Gül Varol, and Andrew Zisserman. Automatic dense annotation of large-vocabulary sign language videos. In *ECCV*, 2022.
- [35] Carol Padden and Darline Clark Gunsauls. How the alphabet came to be used in a sign language. *Sign Language Studies*, 2003.
- [36] Tomas Pfister, James Charles, and Andrew Zisserman. Domain-adaptive discriminative one-shot learning of gestures. In *European Conference on Computer Vision*, pages 814–829. Springer, 2014.
- [37] KR Prajwal, Liliane Momeni, Triantafyllos Afouras, and Andrew Zisserman. Visual keyword spotting with attention. *arXiv preprint arXiv:2110.15957*, 2021.
- [38] Nicolas Pugeault and Richard Bowden. Spelling it out: Real-time asl fingerspelling recognition. In *2011 IEEE International conference on computer vision workshops (ICCV workshops)*, pages 1114–1119. IEEE, 2011.
- [39] Susanna Ricco and Carlo Tomasi. Fingerspelling recognition through classification of letter-to-letter transitions. In *Asian conference on computer vision*, pages 214–225. Springer, 2009.
- [40] Adam Schembri and Trevor Johnston. Sociolinguistic variation in the use of fingerspelling in australian sign language: A pilot study. *Sign Language Studies*, 7(3):319–347, 2007.
- [41] Bowen Shi and Karen Livescu. Multitask training with unlabeled data for end-to-end sign language fingerspelling recognition. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 389–396. IEEE, 2017.

- [42] Bowen Shi, Aurora Martinez Del Rio, Jonathan Keane, Jonathan Michaux, Diane Brentari, Greg Shakhnarovich, and Karen Livescu. American sign language fingerspelling recognition in the wild. In *2018 IEEE Spoken Language Technology Workshop (SLT)*, pages 145–152. IEEE, 2018.
- [43] Bowen Shi, Aurora Martinez Del Rio, Jonathan Keane, Diane Brentari, Greg Shakhnarovich, and Karen Livescu. Fingerspelling recognition in the wild with iterative visual attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5400–5409, 2019.
- [44] Bowen Shi, Diane Brentari, Greg Shakhnarovich, and Karen Livescu. Fingerspelling detection in american sign language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4166–4175, 2021.
- [45] Bowen Shi, Diane Brentari, Greg Shakhnarovich, and Karen Livescu. Searching for finger-spelled content in american sign language. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1699–1712, 2022.
- [46] Themis Stafylakis and Georgios Tzimiropoulos. Zero-shot keyword spotting for visual speech recognition in-the-wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 513–529, 2018.
- [47] Thad 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.
- [48] Rachel Sutton-Spence. Mouthings and simultaneity in british sign language. *Amsterdam studies in the theory and history of linguistic science, series 4*, 281:147, 2007.
- [49] Rachel Sutton-Spence, Bencie Woll, and Lorna Allsop. Variation and recent change in fingerspelling in british sign language. *Language Variation and Change*, 2(3):313–330, 1990.
- [50] Shinichi Tamura and Shingo Kawasaki. Recognition of sign language motion images. *Pattern recognition*, 21(4):343–353, 1988.
- [51] Gul Varol, Liliane Momeni, Samuel Albanie, Triantafyllos Afouras, and Andrew Zisserman. Read and attend: Temporal localisation in sign language videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16857–16866, 2021.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [53] Christian Vogler and Dimitris Metaxas. A framework for recognizing the simultaneous aspects of american sign language. *Computer Vision and Image Understanding*, 81(3):358–384, 2001.
- [54] Christian Vogler and Dimitris Metaxas. Handshapes and movements: Multiple-channel american sign language recognition. In *International Gesture Workshop*, pages 247–258. Springer, 2003.
- [55] Ming-Hsuan Yang, Narendra Ahuja, and Mark Tabb. Extraction of 2d motion trajectories and its application to hand gesture recognition. *IEEE Transactions on pattern analysis and machine intelligence*, 24(8):1061–1074, 2002.

- [56] Polina Yanovich, Carol Neidle, and Dimitris Metaxas. Detection of major asl sign types in continuous signing for asl recognition. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3067–3073, 2016.