

This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Sign Segmentation with Changepoint-Modulated Pseudo-Labelling

Katrin Renz^{1,2} Nicolaj C. Stache² Neil Fox³ Gül Varol^{1,4}

Samuel Albanie¹

² University of Heilbronn, Germany
³ Deafness, Cognition and Language Research Centre, University College London, UK
⁴ LIGM, École des Ponts, Univ Gustave Eiffel, CNRS, France

¹Visual Geometry Group, University of Oxford, UK

https://www.robots.ox.ac.uk/~vgg/research/signsegmentation

Abstract

The objective of this work is to find temporal boundaries between signs in continuous sign language. Motivated by the paucity of annotation available for this task, we propose a simple yet effective algorithm to improve segmentation performance on unlabelled signing footage from a domain of interest. We make the following contributions: (1) We motivate and introduce the task of source-free domain adaptation for sign language segmentation, in which labelled source data is available for an initial training phase, but is not available during adaptation. (2) We propose the Changepoint-Modulated Pseudo-Labelling (CMPL) algorithm to leverage cues from abrupt changes in motionsensitive feature space to improve pseudo-labelling quality for adaptation. (3) We showcase the effectiveness of our approach for category-agnostic sign segmentation, transferring from the BSLCORPUS to the BSL-1K and RWTH-PHOENIX-Weather 2014 datasets, where we outperform the prior state of the art.

1. Introduction

Sign languages are visuo-gestural, evolved languages that represent the natural means of communication for deaf communities [56]. Automatic systems for recognising and understanding signing content have a wide range of applications: enabling indexing of signing content to facilitate efficient search, assistive tools for education and sign linguistics analysis, and sign "wake-word" recognition for virtual assistants [5, 48].

A major challenge in developing such systems is the relative paucity of annotated sign language data that may be employed for training [5, 32]. Several factors drive this state of affairs: a limited supply of annotators with requisite knowledge of sign language required to perform labelling and the extremely high cost of producing the labels themselves [17].

In this work, we focus our attention on the task of sign



Figure 1. Changepoint-Modulated Pseudo-Labelling for sign segmentation: Following an initial pretraining phase on annotated data (top), our proposed approach leverages changepoints detected in motion-sensitive feature space to enhance the quality of annotation produced by pseudo-labelling video data on the target domain of interest (bottom) to perform *source-free adaptation* (target domain adaptation without concurrent access to the source domain). Sign boundaries are marked with blue bars.

segmentation and propose an automatic method for temporally localising sign boundaries in videos of continuous sign language. Sign segmentation plays an important role in the construction of sign language corpora [13, 14] and therefore the development of tools that can perform this task automatically offer the potential to alleviate the limited supply of labelled corpora currently available.

A number of factors make the objective of localising boundaries in continuous sign language challenging. To solve the task, a model must be capable of discriminating gestural patterns from hands moving at high speed, typically in the presence of motion blur, as well as subtle cues in other modalities such as facial expressions. Moreover, in contrast to problem settings such as human action segmentation which have profited from a diverse collection of densely annotated datasets [16, 20, 22, 36, 55, 63], existing sign language datasets provide a relatively small quantity of data with precise temporal boundaries [27, 52], limiting their ability to train models that can generalise robustly beyond the training domain.

An additional factor of particular relevance to sign language datasets relates to the need to protect the privacy of individuals who have contributed human data. In particular, while marginalized communities may potentially benefit from systems (such as those for automatic sign language understanding) that are designed to meet their needs, they face higher risk by contributing data to those systems [6]. This is because smaller group size makes personal identification easier, and marginalized status renders any privacy breaches that occur more dangerous.

In this work, we therefore consider the setting in which the signing data to be segmented, the "target" domain, is sourced from a different distribution to the videos for which segmentation annotation was available for training the model, the "source" domain (see Fig. 4 for examples of different domains considered in this work). We further assume that while the source domain data is available for training an initial model, it is no longer accessible when adapting to the target domain, reflecting the assumption that it may be feasible to share models trained on human data, but infeasible to share the original training data itself.

To make use of the unlabelled data that is available on the target domain, we propose an approach inspired by the classical technique of pseudo-labelling, in which a classifier is retrained on its own predictions on unlabelled data to improve performance [39]. This method, which has proven popular in the context of semi-supervised learning, leverages implicitly the *cluster assumption*—namely that the decision boundaries of the model should lie in regions of low density [10] to adapt the model to new examples. While this assumption is reasonable in the absence of other knowledge, it can be problematic when transferring a sign segmentation model from one domain to another when invariances learned on the former are inappropriate on the latter. Particularly for a small source training set, there may be forms of sign boundaries to which the segmentation model has simply not been exposed, and the discovery of such examples through entropy minimisation schemes [24] such as pseudolabelling is unlikely.

Consistent with this hypothesis, we show through experiments in Sec. 4 that pseudo-labelling provides a boost to performance but exhibits a subtle systematic bias towards under-segmentation. To address this issue, one finding of this work is that it is possible to correct this bias by encoding a simple assumption into the pseudo-labelling process, namely that *sign boundaries typically correspond to motion disfluencies*. To this end, we propose an extension to pseudo-labelling that uses changepoint detection among sequences of motion-sensitive features to modulate the labelset produced on the target domain, increasing their sensitivity to abrupt changes in feature space (see Fig. 1).

We make three contributions: (1) We motivate and introduce the task *source-free* domain adaptation for sign language segmentation, in which labelled source data is available for an initial training phase, but is not available during adaptation. To the best of our knowledge, despite its importance, this task has not been previously investigated in the literature. (2) We establish baselines for this task and propose a modification to pseudo-labelling, which we term *Changepoint-Modulated Pseudo-Labelling*, to address the under-segmentation bias exhibited by naive pseudolabelling on the target domain. (3) We showcase the effectiveness of our approach for category-agnostic sign segmentation, transferring from the BSLCORPUS [52, 53] to the BSL-1K [1] and RWTH-PHOENIX-Weather 2014 [33] datasets, where we outperform the prior state of the art.

2. Related Work

Our work relates to several themes that have been explored in the literature: *temporal action segmentation*, *pseudo-labelling techniques*, *source-free domain adaptation*, *changepoint detection* and *sign language segmentation*.

Action segmentation. The temporal segmentation of untrimmed videos into sequences of actions has received a great deal of attention in the literature, leading to the development of temporal sliding window classifiers [29, 49], generative techniques with Hidden Markov Models (HMMs) [37, 57] and stochastic grammars [59]. More recently, the considerable effectiveness of temporal convolutional networks has been demonstrated for action segmentation, notably with the introduction of Multi-Stage Temporal Convolutional Networks [19] (which we employ in our approach). Of relevance to our approach, the work of [12] explored the task of domain adaptation for action segmentation networks. While effective, their approach requires continued access to the source domain to enable feature alignment between domains (in common with many other unsupervised domain adaptation methods [2, 21, 46, 60]), and is therefore not applicable in our problem formulation.

Pseudo-labelling. The use of pseudo-labelling schemes for exploiting unlabelled data to improve performance has a long history of study stretching back to the 1960s [54, 62]. Variants of this idea have been explored in semi-supervised learning, in which the model is assumed to have access to both labelled and unlabelled data from which to learn from [10, 24]. One particular formulation of this idea known as *pseudo-labelling* [39] has emerged as an especially effective mechanism for semi-supervised learning in which the predictions of a classification model are discretized into one-hot categorical labels and assigned to unlabelled examples, then mixed in with labelled examples

to provide an updated training set for the model. A form of pseudo-labelling was also recently considered for the task of sign recognition by [40] who proposed to localise additional training samples in unlabelled news footage to achieve greater robustness. We employ pseudo-labelling in the specialised *source-free* setting in which no labelled examples are available after the initial training phase on the source domain, discussed next.

Source-free domain adaptation. The adaptation of a model to unlabelled data on a target domain of interest without concurrent access to labelled source domain data becomes more important due to privacy policies in certain domains. Recently the task has been explored using generative models [26, 38, 41], with class prototypes [61] or pseudo-labelling [43]. In this work, we also use pseudo-labelling to adapt models but in addition, we seek to explicitly encode knowledge about the sign segmentation task through changepoint detection.

Changepoint detection. Our category-agnostic sign boundary detection naturally relates to changepoint detection algorithms [8, 11, 45], in which the goal is to locate state changes in time series in order to segment the underlying signal. We refer the reader to [58] for a detailed overview of offline changepoint detection methods. In this work, we propose a learning framework in which we integrate the bottom-up changepoint detections in our highdimensional video features obtained from motion-sensitive pretraining tasks.

The segmentation of continuous Sign segmentation. streams of signing into individual sign "tokens" has been the subject of considerable interest in the sign linguistics community, with a particular focus on how to define appropriate boundaries [14, 15, 25] and measuring consistency between annotation teams [4, 23]. By contrast, automatic sign segmentation has received relatively limited attention in the computational literature. Of the prior research in this area, the dominant approach has been to employ methods that require a semantic labelling of the signing content (this can take the form of glosses-the minimal lexical units of signing used for annotation, or full translations) [35, 51]. This differs from continuous sign language recognition (CSLR) [34] which aims to determine sign order without localising boundaries (note that while such methods may produce sign boundaries implicitly, they likewise require access to dense semantic labels). Other methods have proposed to tackle related tasks such as identifying whether a person is actively signing [3, 44] and parsing continuous sign language into sentence-like units [7]. Of most relevance to our work, the recent works of [18] and [47] proposed methods to tackle the category-agnostic segmentation problem. In [18], the authors investigated the use of a random forest with geometric features derived from 3D skeleton data (gathered via motion capture). The later work of [47] demonstrated the superiority of the MS-TCN

architecture [19] for this task, and we therefore adopt this model in our approach. Differently from [47] who assume the training and test data are drawn from the same distribution, we consider the setting in which the target domain differs from source domain.

3. Method

We propose a simple adaptation method for improving sign segmentation performance on a target domain for which labelled data is unavailable under the assumption that the source domain (containing labelled data) may not be accessed during the adaptation process (we formalise this problem definition in Sec. 3.1). Our approach employs a standard fully-supervised training phase on the source domain (Sec. 3.2), before running a secondary phase of iterative adaptation with the *Changepoint-Modulated Pseudo-Labelling* algorithm (described in Sec 3.3).

3.1. Problem formulation

Let \mathcal{X} denote the set of videos containing a person performing continuous sign language. The goal of temporal sign segmentation is to produce for each $\mathbf{x} \in \mathcal{X}$ with frames (x_1, \ldots, x_N) a corresponding set of frame labels $\mathbf{y} = (y_1, \dots, y_N) \in \{0, 1\}^N$, in which label values of 1 and 0 denote boundaries between sign segments and the interior of segments, respectively. We assume access to a set of labelled sign segments that are gathered from a "source" domain $\mathcal{X}_S \subset \mathcal{X}$ on which an initial model may be trained to perform the segmentation task. We further assume that this domain differs from that of the "target" videos $\mathcal{X}_T \subset \mathcal{X}$ of interest, for which annotations are not available. The objective of this work is to maximise sign segmentation performance on \mathcal{X}_T , under the constraint that at no point we have concurrent access to video samples from both the source and target domain.

3.2. Source domain training

To make best use of labelled data in the source domain, we first train a sign segmentation model in a fullysupervised manner. Following the dominant approaches in the temporal segmentation literature [19, 42], we assume that the sign segmentation model decomposes into a visual feature extractor, $\phi(\mathbf{x})$, (typically instantiated as a spatiotemporal convolutional network such as I3D [9]) and a segmentation network, $\psi(\phi)$ which ingests these features and outputs per-frame segmentation labels [19, 47]. In this work, we employ a standard frame-level cross-entropy loss in combination with the smoothing loss proposed by [19] to train $f = \psi \circ \phi$. Once source domain training has completed, the source videos and labels are discarded and only the trained model, f, is retained.



Figure 2. **Method overview:** The proposed approach extracts I3D features and employs both semantic and low-level cues to produce labels on the target dataset (block 1) via the CMPL algorithm. These labels are then employed directly via finetuning (block 2).

3.3. Changepoint-Modulated Pseudo-Labelling

The key idea behind our adaptation approach is to leverage the benefits of *pseudo-labelling* [39] while mitigating its limitations for the task of sign language segmentation through a simple technique we term *changepointmodulation*. A method overview can be seen in Fig. 2.

Pseudo-labelling works by minimising the conditional entropy of a given classification model across unlabelled samples. It achieves this by assigning to each sample the class label corresponding to the maximum posterior probability of the model and then retraining the model to become more confident in its predictions. In doing so, it encodes the intuition that individual samples should not represent a mixture of classes-rather they should belong to one or the other (corresponding to the cluster assumption that decision boundaries should lie in regions of low-density [10]). In this work we employ pseudo-labelling in the manner described above: we perform inference with the sign segmentation model $f = \psi \circ \phi$ (that was trained on the source domain) across videos of unlabelled signing drawn from \mathcal{X}_T to produce frame-level posterior probabilities for each frame. We then assign each frame the binary label corresponding to the maximum posterior probability of the model and retrain the model by employing the same losses that were used on the source domain (Sec. 3.2).

As we show through experiments in Sec. 4 (and consistent with the literature [39, 50, 64]), pseudo-labelling yields a boost in performance on the target domain. However, while it brings consistent improvement, we observe that in practice the algorithm exhibits a subtle but systematic bias towards *under-segmenting* signs. In the absence of other knowledge, the cluster assumption represents a reasonable prior to improve target domain performance. However, for the particular task of sign segmentation, the sign linguistics literature suggests that there are particular gestural movements and cues that correlate with sign boundaries [25, 28] which we can expect to hold true across domains. While prescribing a precise set of rules that define a sign boundary would be extremely challenging [25], we can never-



Figure 3. **CMPL components:** We show the effect of the insertion and refinement pseudo-label transformations on a sequences of pseudo-labels (PL) and changepoints (CP) relative to ground truth (GT) (see Sec. 3.3 for details). (Top) An example where the changepoints help to mitigate under-segmentation in the PL sequence. (Bottom) Small refinements on the pseudo-labels lead to improved boundary locations.

theless integrate our knowledge about the nature of these boundaries—in particular, their weak correspondence with human motion disfluencies—into the pseudo-labelling process. To do so, we propose to couple motion-sensitive features (which may be learned on a host of human activitybased visual understanding tasks) with changepoint detection in feature space (implementation details are given in detail in Sec. 4.3). In this way, we obtain, in addition to the pseudo-labels described above, an additional set of candidate frame-locations at which a sign boundary is more likely to occur (those that correspond to discontinuities in feature space detected by the changepoint algorithm).

A natural question then arises: how should the two sets of candidate frame-level labels be integrated to improve target domain performance? One simple strategy would be to take the union of boundaries predicted by both methods. However, this strategy, as with that of simple averaging, does not account for the fact that intuitively, we would like pseudolabels to be reinforced when they appear in a close neighbourhood of a motion feature changepoint without requiring exact alignment to achieve this effect.

To address this, we propose the *Changepoint-Modulated Pseudo-Labelling* algorithm, which comprises two pseudolabel transformations. Motivated by the under-segmentation issue described above, the first transformation inserts new boundaries suggested by abrupt changes in feature space, when far from pseudo-label boundaries. More concretely, given a sequence of target-domain frame-level binary pseudo-labels, $(\hat{y}_1^T, \ldots, \hat{y}_N^T)$, and frame-level binary pseudo-labels, (c_1^T, \ldots, c_N^T) (obtained by placing a unit value at every frame corresponding to a changepoint), we perform the following pseudo-label *insertion* transformation:

$$\hat{y}_i^T = \begin{cases} c_i^T, & \text{if } \sum_{j \in (-\gamma, \gamma)} \hat{y}_{i+j}^T = 0\\ \hat{y}_i^T, & \text{otherwise} \end{cases}$$
(1)

where γ represents a bandwidth value (set as a hyperparameter). The second *refinement* transformation aims to minimise potential bias towards the annotation style used in the source dataset. First, each contiguous sequence of boundary labels in the pseudo-labels is matched to the nearest contiguous sequence of boundary labels in the changepoint sequence, under the condition that it falls within a matching window of δ frames (also set as a hyperparameter). Matched contiguous pseudo-label boundary sequences are then translated in time such that their central element occupies the midpoint between the original pseudo-label and changepoint boundary positions prior to performing the translation. Examples for these transformation steps can be seen in Fig. 3.

Following these transformations, the sign segmentation model is simply retrained on the updated set of pseudolabels across the target domain videos. We conduct ablations to assess the utility of both the *insertion* and the *refinement* transformations in Sec. 4, as well as the role of the corresponding bandwidth and matching neighbourhood hyperparameters, γ and δ .

4. Experiments

This section describes the datasets (Sec. 4.1), the evaluation metrics (Sec. 4.2), and further implementation details (Sec. 4.3). We then present various baselines, ablations and comparisons to the prior state of the art (Sec 4.4). Finally, we provide qualitative results (Sec 4.5).

4.1. Datasets

We use three datasets in our experiments. Sample video frames from each of these datasets are provided in Fig. 4.

BSLCORPUS [52, 53] is a British Sign Language (BSL) dataset which contains, inter alia, continuous signing in the form of conversations or narratives. For a subset of the data, fine-grained annotations exist which contain start and end times, as well as the sign categories. We use this as the source dataset to perform pretraining. We employ the train and test partitions used in [47] which contain approximately six hours of signing with gloss-level annotations. A detailed description of the dataset can be found via [52].

BSL-1K [1] consists of public broadcast footage of sign language interpreted videos, organised into a number of episodes. This dataset shares the same language as BSLCORPUS (i.e., BSL) but differs in content, background, person position, and signing speed. To evaluate the performance on this dataset we use the manually-annotated subset of signing data provided by [47], comprising two minutes of video footage with 177 sign instances. We experiment



Figure 4. Sample frames from the three video datasets considered in this work: We investigate source-free domain adaptation for sign segmentation across different domains of discourse (personal narratives vs TV broadcast content), context (conversational vs interpreted signing) and sign language (BSL vs DGS).

with training on a subset of three episodes of BSL-1K corresponding to 3.5 h of signing content. We also conduct experiments under a *transductive setting* in which we perform adaptation directly on the episode from which the labelled test set sequence was sourced (this episode represents 45 min of signing content).

RWTH-PHOENIX-Weather 2014 [33] (PHOENIX14) is a German Sign Language (DGS) dataset that has been widely employed for the study of automatic sign language recognition. The work of [35] provides automatic temporal segment labels for the training partition of this data, which they obtain through forced-alignment (using ground-truth gloss information). We use these labels against which we evaluate our predictions. To enable direct comparison, we employ the same train/test partitions as used in [47] resulting in a 1.76 h test set.

4.2. Evaluation metrics

We use the same metrics as in [47] and report for each the mean and standard deviation out of three runs with different random seeds. They use two different metrics: mF1B to measure the performance of the boundary position and mF1S for the extent of the sign segments.

mF1B. One boundary, which we define as a series of following 1s in y, is predicted correct, if the distance to a ground truth boundary is lower than a certain threshold. We use the mean of all F1 scores with thresholds in the interval [1, 4].

mF1S. A sign segment is counted as correct if the IoU is higher than a given threshold. We calculate the average F1 score for all thresholds in the range of 0.4 and 0.75 with the step size 0.05. This metric is, in contrast to the mF1B metric, sensitive to the width of the predicted and ground truth boundaries, which leads also to sensitivity to different anno-

| | Adaptation protocol | mF1B | mF1S |
|---------------------------------------|--|--|--|
| Source-only [47] | | $46.75_{\pm 1.2}$ | $32.29_{\pm 0.3}$ |
| Pseudo-labels Changepoints CMPL | inductive inductive inductive | $\begin{array}{c} 47.94_{\pm 1.0} \\ 48.51_{\pm 0.4} \\ \textbf{53.57}_{\pm \textbf{0.7}} \end{array}$ | $\begin{array}{c} 32.45_{\pm 0.3} \\ 34.45_{\pm 1.4} \\ 33.82_{\pm 0.0} \end{array}$ |
| Pseudo-labels Changepoints CMPL | transductive transductive transductive | $\begin{array}{c} 47.62_{\pm 0.4} \\ 48.29_{\pm 0.1} \\ \textbf{53.53}_{\pm \textbf{0.1}} \end{array}$ | $\begin{array}{c} 32.11_{\pm 0.9} \\ 35.31_{\pm 1.4} \\ 32.93_{\pm 0.9} \end{array}$ |

Table 1. **Results on the PHOENIX14 dataset:** (Top) Naive sign segmentation performance without assuming any target domain labels. (Middle and Bottom) A comparison of adaptation strategies in the inductive (middle) and transductive (bottom) setting. In each case, we see that the proposed Changepoint-Modulated Pseudo-Labelling method (CMPL) outperforms alternatives.

tation styles between annotators. To avoid this, we consider the mF1B metric as our main metric for ranking the results but include mF1S for completeness.

4.3. Implementation details

Features. For all datasets, we employ an I3D [9] backbone architecture, ϕ , pretrained for the task of action recognition on the Kinetics dataset [30] to provide representations that are sensitive to fine-grained human motions. To ensure viable adaptation on any unlabelled data, we do not use any annotations from the target domain; therefore we train the model directly on the class labels of BSLCORPUS.

MS-TCN pretraining on the source domain. For the source domain training phase described in Sec. 3.2, we pretrain an MS-TCN model [19] as our sign segmentation network, $f = \psi \circ \phi$, that further processes the outputs of the I3D feature extractor ϕ described above with a segmentation network ψ comprising a stack of temporal convolutional layers. The model is trained under a per-frame binary-class classification objective. Similar to [19], we use an architectural design of 4 stages with 10 layers in each stage, 64 filters, and an Adam optimizer.

Changepoint-Modulated Pseudo-Labelling implementation details and hyperparameters. We adopt the Pelt [31] method, an exact changepoint detection algorithm which exhibits linear average runtime behaviour, as a basis for discovering disfluencies in signing from the space of I3D features. We employ the implementation provided by [58] with an L2 cost function on I3D features exacted with a stride of one frame. We use a Pelt cost penalty term of 100 when using changepoints to modulate pseudo-label insertion and refinement. Both the bandwidth and matching window hyperparameters (γ and δ) introduced in Sec. 3.3 are set to a value of four frames. Since boundaries estimated by the changepoint algorithm span only a single frame, we expand the changepoints to span three frames in terms of sign boundary width (all videos considered in this work are encoded at 25 fps). This ensures that the boundary estimations are similar in width to the pseudo-labels produced by the MS-TCN sign segmentation network.

| PL threshold | mF1B | mF1S |
|--------------|-------------------|-------------------|
| 0.5 | $47.94_{\pm 1.0}$ | $32.45_{\pm 0.3}$ |
| 0.4 | $48.01_{\pm 1.0}$ | $31.26_{\pm 0.7}$ |
| 0.3 | $48.21_{\pm 1.0}$ | $29.87_{\pm 1.1}$ |
| 0.2 | $46.67_{\pm 0.6}$ | $27.52_{\pm 0.7}$ |

Table 2. **Pseudo-label training with reduced thresholds on the PHOENIX14 dataset:** Reducing the threshold at which a pseudolabel probability is mapped to a boundary label does not address the under-segmentation issue in a way that improves performance.

| | Adaptation protocol | mF1B | mF1S |
|----------------------------------|---------------------|-------------------|-------------------|
| Baselines: | | | |
| Uniform (using GT #signs) | - | 41.80 | 34.75 |
| Changepoints (using GT #signs) | - | 60.73 | 52.89 |
| Changepoints (estimating #signs) | - | 60.25 | 53.94 |
| Prior works: | | | |
| Geometric features + RF [18, 47] | source-only | $51.26_{\pm 0.5}$ | $34.28_{\pm 1.0}$ |
| MS-TCN [19, 47] | source-only | $61.12_{\pm 0.9}$ | $49.96_{\pm 0.6}$ |
| Proposed model: | | | |
| CMPL | inductive | $65.99_{\pm 1.0}$ | $48.81_{\pm 1.3}$ |
| CMPL | transductive | $67.01_{\pm 2.2}$ | $50.20_{\pm 0.6}$ |

Table 3. **Results on the BSL-1K dataset:** We compare our approach to several baselines and the previous state of the art. The changepoint-only approach already builds a strong baseline, especially for the mF1S metric. With CMPL we gain about 6% mF1B over the naive transfer from BSLCORPUS-labelled training of [47].

Training and evaluation. To evaluate performance, we report the improvement derived from performing adaptation using unlabelled training set videos from the target domain (*inductive* setting). Since our approach does not require the use of labels on the target dataset, we also report results for experiments in which perform adaptation directly on the test data without labels (*transductive setting*), a formulation which can often arise in practice when the target videos to be segmented are known in advance. Due to the small portion of BSL-1K with annotations available, we do not use a separate validation set, and fix the number of epochs to 10.

4.4. Ablation studies

PHOENIX14. To assess the influence of components of our framework, we first present an ablation study on the PHOENIX14 dataset in Tab. 1, in which we compare against the baseline performance of [47] trained only on source data. By using the pseudo-labelling technique, we obtain a first small improvement (46.75 vs 47.94 mF1B). It is interesting to note that the changepoint detection algorithm alone establishes a very strong baseline (48.51 mF1B). More generally we observe that self-training with both pseudo-labels and changepoints provide a significant boost in adapting the model to the new domain under both inductive and transductive model evaluations. However, we see that our Changepoint-Modulated Pseudo-Labelling (CMPL) provides a significantly greater boost in mF1B (of about 6 points).

We also show that the under-segmentation issue is not addressed by simply reducing the threshold employed to con-



Figure 5. Sensitivity to the threshold hyperparameters of the CMPL method for BSL-1K (transductive): We show the influence of different values for thresholds for the *insertion* bandwidth γ (left) and *refinement* matching window δ (right).

vert pseudo-label posterior probabilities into boundary class labels (for all other experiments we simply convert each frame that is assigned a boundary probability higher than 0.5 to a boundary label). Tab. 2 reports the performance of the model under the influence of varying this threshold. See Sec. A of the appendix for further ablations on PHOENIX14.

BSL-1K. We observe that in both the inductive and transductive setting, CMPL yields a significant gain over the strongest reported result for sign segmentation on BSL-1K. In particular, we observe an improvement of about 6% mF1B over the naive transfer of the model trained on the source data. To better understand the behaviour of the proposed approach, we next conduct further ablations using various baselines on the BSL-1K dataset. We describe these baselines next and report their performance in Tab. 3. We also show the performance of alternative fusion strategies and the hyperparamter sensitivity of our approach.

Uniform baseline. We report the performance of uniformly splitting the target video into segments of equal size such that the total number of segments match the total number of ground-truth boundaries. Since the uniform baseline assumes that the ground-truth number of signs are known, it is thus *not* directly comparable to the automatic segmentation models, but provides a degree of calibration for the difficulty of the task.

Changepoints-only baseline. For the changepoint baseline, also reported in Tab. 3, we employ only the feature extraction network, ϕ , and calculate the changepoints on the extracted features, $\phi(\mathbf{x})$. The first method uses dynamic programming with the ground-truth number of boundaries given. By using the changepoint algorithm, we observe a large improvement in both metrics over uniformly splitting the video. However, we also observe that by using the Pelt method to predict position and number of boundaries we come close to matching the variant that uses ground truth information, suggesting the robustness of this technique and its suitability as a building block in our approach. In Fig. 9, we assess the sensitivity of this approach to the choice of different values for the penalty hyperparameter and determine that a choice of 80 represents a reasonable trade-off



Figure 6. Sensitivity to the penalty hyperparameter of the CMPL method for BSL-1K (transductive): We show the influence of different values for varying the changepoint penalty when fixing the threshold hyperparameters to $\gamma = 4$, $\delta = 4$ as used in all other experiments.

| Fusion strategies | Adaptation protocol | mF1B | mF1S |
|--|------------------------------|--|---|
| merge PL + CP local fusion | inductive inductive | ${}^{65.10_{\pm 1.6}}_{62.35_{\pm 1.4}}$ | ${}^{42.73_{\pm 1.0}}_{48.77_{\pm 2.3}}$ |
| insertion insertion + refinement (CMPL) | inductive inductive | $\begin{array}{c} 62.49_{\pm 2.2} \\ \textbf{65.99}_{\pm 1.0} \end{array}$ | $\begin{array}{c} 45.77_{\pm 1.6} \\ 48.81_{\pm 1.3} \end{array}$ |
| merge PL + CP local fusion | transductive transductive | $\begin{array}{c} 65.69_{\pm 1.3} \\ 62.71_{\pm 3.1} \end{array}$ | $\begin{array}{c} 43.06_{\pm 3.0} \\ 50.69_{\pm 1.9} \end{array}$ |
| insertion insertion + refinement (CMPL) | transductive transductive | $\begin{array}{c} 63.27_{\pm 3.4} \\ 67.01_{\pm 2.2} \end{array}$ | $\begin{array}{c} 48.49_{\pm 3.1} \\ 50.20_{\pm 0.6} \end{array}$ |

Table 4. **Fusion strategies:** We compare several strategies for fusing pseudo-labels and changepoints on the BSL-1K test set. We observe that the combined "insertion + refinement" strategy proposed as part of the CMPL approach consistently performs best.

for the changepoints-only baseline.

Effect of the fusion strategy. We explore several alternative ways to fuse the two knowledge sources provided by pseudo-labels and changepoint detections. The simplest such strategy takes the union of predictions and keeps all boundaries which are present either in the pseudo-labels or in the changepoints (labelled "merge PL + CP" in Tab. 4). We also consider an alternative strategy that specifically targets the under-segmentation problem induced by a pseudolabel-only approach. To insert more boundaries we select those segments which are longer than the average sign length and insert detected changepoints only amongst such regions. These results are reported in Tab. 4 under the row title "local fusion". Finally, we evaluate our proposed CMPL fusion strategy and investigate the impact of the two different pseudo-label transformations that it employs. In Tab. 4, we observe when used in isolation the *insertion* transformations produce some improvement, but that there is a clear benefit to the combination of insertion and refinement, justifying the usage of this combination in our design. Hyperparameter sensitivity. We further show the sensitivity to the hyperparameters of the CMPL method. First we show the impact of different values for thresholds for the insertion bandwidth γ and refinement matching window δ in Fig. 5. Changes to these hyperparameters produce a small variation in the results. We observe that an insertion bandwidth of $\gamma = 4$ frames and a refinement matching window of $\delta = 4$ frames with the given changepoint penalty of 100



Figure 7. **Qualitative results on PHOENIX14 [33]:** We compare the results of the different models, source-only (SO), pseudolabels (PL), changepoints (CP), Changepoint-Modulated Pseudo-Labelling (CMPL) with the ground truth (GT). We show two different extracts, in which our method is able to detect new and suppress wrong boundaries (green box). The red box indicates failure cases, where further work is necessary.

works well. Next, in Fig. 6, we fix $\gamma = 4$ and $\delta = 4$ and change the changepoint penalty. We observe the best result in the mF1B metric for a penalty of 100. Looking at the mF1S metric, leads to the assumption, that this metric gains from the insertion of fewer boundaries.

4.5. Qualitative analysis

In Fig. 3, we illustrate two examples of the proposed Changepoint-Modulated Pseudo-Labelling method in action. The upper one depicts an extract of an episode in which the described under-segmentation problem of the pseudo-labels (PL) can be clearly observed. As the changepoint detection finds a boundary and no pseudo-label is in the neighbouring area, we insert a new boundary at the detected position, remedying the omission. In the lower example, we observe the behaviour of the refinement transformation and its role in improving overall pseudo-label quality by adjusting the boundary locations to account for evidence from the changepoint detections.

In Fig. 7 and Fig. 8 we provide a qualitative comparison of the different models on two sequences from the PHOENIX14 and BSL-1K datasets. We show failure cases in red and success cases in green. In Fig. 7, the upper example demonstrates the effectiveness of the CMPL for predicting new boundaries through the changepoint insertion. The first two false positive boundaries in the bottom example are due to a repeated sign, which leads to abrupt changes at the time the sign gets repeated. The last boundary in the red frame, which is falsely detected by each method, is aligned with a strong direction change in the signing. In general, we observed in our error analysis two main issues: (i) the new method mainly helps against under-segmentation, but can only in seldom cases suppress false positive predictions in the pseudo-labels or changepoints, (ii) the insertion of the changepoints leads to higher true positive rate, but also to a higher false positive rate. This can be observed especially



Figure 8. **Qualitative results on BSL-1K** [1]: (Top) Changepoints help to find boundaries which were previously missed. (Bottom) Through the insertion of the changepoint the CMPL method detects new boundaries, but is in most cases not able to suppress the false positives detected by the source-only approach.



Figure 9. **Pelt changepoint penalty term:** We study the influence of the penalty value for Pelt changepoint detection algorithm when using only changepoints as labels to train on the target domain. We observe that high penalty results in under segmentation, whereas a low penalty score over-segments.

for PHOENIX14. The previously mentioned two issues can be also observed for BSL-1K (Fig. 8). We detect more correct boundaries (green box), but also some false positives (left red box) and are not able to suppress already existing false positives (right red box). For further qualitative results, we refer to our supplemental video on the project webpage and Sec. B of the appendix.

5. Conclusion

We presented an approach to temporally segment signs in continuous sign language videos, with a particular emphasis on leveraging unlabelled data for training. We motivated and introduced the problem of *source-free* domain adaptation in this context, in which we assume access to labelled source data during model training but not during adaptation to the target domain. We proposed the conceptually simple but powerful Changepoint-Modulated Pseudo-Labelling algorithm and demonstrated its effectiveness through state-of-the-art performance on two sign language datasets. Potential future directions include repurposing our segmentation models for improving sign language recognition and active signer detection performance.

Acknowledgements. This work was supported by EPSRC grant ExTol. KR was supported by the German Academic Scholarship Foundation. The authors thank Andrew Zisserman for suggestions and Cihan Camgöz for assistance with data preparation.

References

- Samuel Albanie, Gül Varol, Liliane Momeni, Triantafyllos Afouras, Joon Son Chung, Neil Fox, and Andrew Zisserman. BSL-1K: Scaling up co-articulated sign language recognition using mouthing cues. In *ECCV*, 2020.
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [3] Mark Borg and Kenneth P Camilleri. Sign language detection "in the wild" with recurrent neural networks. In *ICASSP*, 2019.
- [4] Annelies Braffort and Leïla Boutora. Défi d'annotation degels2012: la segmentation (degels2012 annotation challenge: Segmentation)[in french]. In JEP-TALN-RECITAL 2012, Workshop DEGELS 2012: Défi GEste Langue des Signes (DEGELS 2012: Gestures and Sign Language Challenge), pages 1–8, 2012.
- [5] Danielle Bragg, Oscar Koller, Mary Bellard, Larwan Berke, Patrick Boudreault, Annelies Braffort, Naomi Caselli, Matt Huenerfauth, Hernisa Kacorri, Tessa Verhoef, et al. Sign language recognition, generation, and translation: An interdisciplinary perspective. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, pages 16–31, 2019.
- [6] Danielle Bragg, Oscar Koller, Naomi Caselli, and William Thies. Exploring collection of sign language datasets: Privacy, participation, and model performance. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–14, 2020.
- [7] Hannah Bull, Michèle Gouiffès, and Annelies Braffort. Automatic segmentation of sign language into subtitle-units. In ECCVW, Sign Language Recognition, Translation and Production (SLRTP), 2020.
- [8] Edward Carlstein, Hans-Georg Müller, and David Siegmund. Nonparametric methods in change-point problems: a general approach and some concrete algorithms, volume Volume 23 of Lecture Notes-Monograph Series, pages 99–107. Institute of Mathematical Statistics, 1994.
- [9] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [10] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- [11] Jie Chen and Arjun K Gupta. *Parametric Statistical Change Point Analysis: With Applications to Genetics, Medicine, and Finance; 2nd ed.* Springer, Boston, 2012.
- [12] Min-Hung Chen, Baopu Li, Yingze Bao, Ghassan Al-Regib, and Zsolt Kira. Action segmentation with joint selfsupervised temporal domain adaptation. In *CVPR*, 2020.
- [13] K Cormier, OA Crasborn, and R Bank. Digging into signs: Emerging annotation standards for sign language corpora. In Efihimiou, E.; et al.(ed.), Proceedings of the 7th Workshop on the Representation and Processing of Sign Languages: Corpus Mining, pages 35–40. Portorož, Slovenia: ELRA, 2016.
- [14] Kearsy Cormier and Jordan Fenlon. BSL corpus annotation guidelines. 2014.

- [15] OA Crasborn and IEP Zwitserlood. The corpus ngt: an online corpus for professionals and laymen. In Proceedings of the 3rd Workshop on the Representation and Processing of Sign Languages: Construction and Exploitation of Sign Language Corpora, pages 44–49. Paris: ELRA, 2008.
- [16] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The dataset. In ECCV, 2018.
- [17] P. Dreuw and H. Ney. Towards automatic sign language annotation for the ELAN tool. In *LREC Workshop on the Representation and Processing of Sign Languages: Construction and Exploitation of Sign Language Corpora*, 2008.
- [18] Iva Farag and Heike Brock. Learning motion disfluencies for automatic sign language segmentation. In *ICASSP*, 2019.
- [19] Yazan Abu Farha and Jurgen Gall. MS-TCN: Multi-stage temporal convolutional network for action segmentation. In *CVPR*, 2019.
- [20] Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activities. In CVPR, 2011.
- [21] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- [22] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *ICCV*, 2017.
- [23] Matilde Gonzalez Preciado. Computer vision methods for unconstrained gesture recognition in the context of sign language annotation. PhD thesis, Université de Toulouse, Université Toulouse III-Paul Sabatier, 2012.
- [24] Yves Grandvalet and Yoshua Bengio. Entropy regularization., 2006.
- [25] Thomas Hanke, Silke Matthes, Anja Regen, and Satu Worseck. Where does a sign start and end? segmentation of continuous signing. In *LREC Workshop on the Representation and Processing of Sign Languages: Interactions between Corpus and Lexicon*, 2012.
- [26] Yunzhong Hou and Liang Zheng. Source free domain adaptation with image translation. arXiv:2008.07514, 2020.
- [27] Elena Jahn, Reiner Konrad, Gabriele Langer, Sven Wagner, and Thomas Hanke. Publishing dgs corpus data: Different formats for different needs. In *Proceedings of the Workshop* on the Representation and Processing of Sign Languages at LREC, 2018.
- [28] Robert E Johnson and Scott K Liddell. A segmental framework for representing signs phonetically. *Sign Language Studies*, 11(3):408–463, 2011.
- [29] Svebor Karaman, Lorenzo Seidenari, and Alberto Del Bimbo. Fast saliency based pooling of fisher encoded dense trajectories. In ECCV THUMOS Workshop, 2014.
- [30] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. arXiv:1705.06950, 2017.
- [31] Rebecca Killick, Paul Fearnhead, and Idris A Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*,

107(500):1590-1598, 2012.

- [32] Oscar Koller. Quantitative survey of the state of the art in sign language recognition. *arXiv:2008.09918*, 2020.
- [33] Oscar Koller, Jens Forster, and Hermann Ney. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141:108–125, 2015.
- [34] Oscar Koller, Jens Forster, and Hermann Ney. Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers. *Computer Vision and Image Understanding*, 141:108–125, 2015.
- [35] Oscar Koller, Sepehr Zargaran, and Hermann Ney. Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent CNN-HMMs. In *CVPR*, 2017.
- [36] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goaldirected human activities. In *CVPR*, 2014.
- [37] Hilde Kuehne, Juergen Gall, and Thomas Serre. An end-toend generative framework for video segmentation and recognition. In WACV, 2016.
- [38] Jogendra Nath Kundu, Naveen Venkat, R Venkatesh Babu, et al. Universal source-free domain adaptation. In CVPR, 2020.
- [39] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 2, 2013.
- [40] Dongxu Li, Xin Yu, Chenchen Xu, Lars Petersson, and Hongdong Li. Transferring cross-domain knowledge for video sign language recognition. In CVPR, 2020.
- [41] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In CVPR, 2020.
- [42] Shi-Jie Li, Yazan AbuFarha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. MS-TCN++: Multi-stage temporal convolutional network for action segmentation. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 2020.
- [43] Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? source hypothesis transfer for unsupervised domain adaptation. In *ICML*, 2020.
- [44] Amit Moryossef, Ioannis Tsochantaridis, Roee Aharoni, Sarah Ebling, and Srini Narayanan. Real-time sign language detection using human pose estimation. arXiv:2008.04637, 2020.
- [45] E. S. PAGE. A test for a change in a parameter occurring at an unknown point. *Biometrika*, 42(3-4):523–527, 1955.
- [46] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. arXiv:1809.02176, 2018.
- [47] Katrin Renz, Nicolaj Stache, Samuel Albanie, and Gül Varol. Sign language segmentation with temporal convolutional networks. In *ICASSP*, 2021.
- [48] Jason Rodolitz, Evan Gambill, Brittany Willis, Christian Vogler, and Raja Kushalnagar. Accessibility of voiceactivated agents for people who are deaf or hard of hearing. *The Journal on Technology and Persons with Disabilities*, page 144, 2019.
- [49] Marcus Rohrbach, Sikandar Amin, Mykhaylo Andriluka, and Bernt Schiele. A database for fine grained activity detection of cooking activities. In *CVPR*, 2012.
- [50] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada.

Asymmetric tri-training for unsupervised domain adaptation. *arXiv:1702.08400*, 2017.

- [51] Pinar Santemiz, Oya Aran, Murat Saraclar, and Lale Akarun. Automatic sign segmentation from continuous signing via multiple sequence alignment. In *ICCV Workshops*, 2009.
- [52] Adam Schembri, Jordan Fenlon, Ramas Rentelis, and Kearsy Cormier. British Sign Language Corpus Project: A corpus of digital video data and annotations of British Sign Language 2008-2017 (Third Edition), 2017.
- [53] Adam Schembri, Jordan Fenlon, Ramas Rentelis, Sally Reynolds, and Kearsy Cormier. Building the British sign language corpus. *Language Documentation & Conservation*, 7:136–154, 2013.
- [54] H Scudder. Probability of error of some adaptive patternrecognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965.
- [55] Sebastian Stein and Stephen J McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 729–738, 2013.
- [56] Rachel Sutton-Spence and Bencie Woll. *The Linguistics of British Sign Language: An Introduction*. Cambridge University Press, 1999.
- [57] Kevin Tang, Li Fei-Fei, and Daphne Koller. Learning latent temporal structure for complex event detection. In *CVPR*, 2012.
- [58] Charles Truong, Laurent Oudre, and Nicolas Vayatis. Selective review of offline change point detection methods. *Signal Processing*, 167:107299, 2020.
- [59] Nam N Vo and Aaron F Bobick. From stochastic grammar to bayes network: Probabilistic parsing of complex activity. In *CVPR*, 2014.
- [60] Riccardo Volpi, Pietro Morerio, Silvio Savarese, and Vittorio Murino. Adversarial feature augmentation for unsupervised domain adaptation. In *CVPR*, 2018.
- [61] Shiqi Yang, Yaxing Wang, Joost van de Weijer, and Luis Herranz. Unsupervised domain adaptation without source data by casting a bait. *arXiv:2010.12427*, 2020.
- [62] David Yarowsky. Unsupervised word sense disambiguation rivaling supervised methods. In 33rd annual meeting of the association for computational linguistics, pages 189–196, 1995.
- [63] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 2017.
- [64] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In ECCV, 2018.