Towards Zero-shot Sign Language Recognition

Yunus Can Bilge, Ramazan Gokberk Cinbis, and Nazli Ikizler-Cinbis

Abstract—
This paper tackles the problem of zero-shot sign language recognition (ZSSLR), where the goal is to leverage models learned over the seen sign classes to recognize the instances of unseen sign classes. In this context, readily available textual sign descriptions and attributes collected from sign language dictionaries are utilized as semantic class representations for knowledge transfer. For this novel problem setup, we introduce three benchmark datasets with their accompanying textual and attribute descriptions to analyze the problem in detail. Our proposed approach builds spatiotemporal models of body and hand regions. By leveraging the descriptive text and attribute embeddings along with these visual representations within a zero-shot learning framework, we show that textual and attribute based class definitions can provide effective knowledge for the recognition of previously unseen sign classes. We additionally introduce techniques to analyze the influence of binary attributes in correct and incorrect zero-shot predictions. We anticipate that the introduced approaches and the accompanying datasets will provide a basis for further exploration of zero-shot learning in sign language recognition.

Index Terms—Sign language recognition, zero-shot learning.

arXiv:2201.05914v1 [cs.CV] 15 Jan 2022

1 Introduction

SIGN language recognition (SLR) is among the open problems with great practical importance in computer vision. One of the factors that render SLR a unique problem is the fact that although most signs have clear-cut definitions, most of them have only subtle visual differences across them [1], [2]. Therefore, towards turning SLR into a ubiquitous technology, fundamental progress in modeling and recognizing fine-grained spatiotemporal patterns expressed by the movement of hands with various shapes, orientations, and locations, as well as body posture, and non-manual features, such as facial expressions, are needed. Besides, the photometric and geometric factors, such as viewpoint changes [3], variations in sign languages across regions [4], and the fact that sign languages embrace significant variations over time [5] render the task more challenging.

The existing SLR approaches typically require a large number of annotated examples for each sign class of interest [6], [7], [8], [9], [10]. This dependency means that one presumably needs to collect annotated samples for all signs in all sign languages of interests, including variations expressed by multiple persons per sign under various recording conditions. Noting that over 140 sign languages and many more dialects are estimated to exist around the world [11], with typically 2500 to 5500 signs per language [12], [13], [14], the need for supervised examples create a *data bottleneck* problem in scaling up SLR. Towards bypassing this problem, we explore the possibility of recognizing sign classes with no annotated visual examples, purely based on existing definitions, which we call *zero-shot sign language recognition* (ZSSLR). Unlike the traditional supervised SLR, where training data is needed for all classes of interest, in ZSSLR, the aim is to recognize novel sign classes that do not have any visual training samples. We additionally introduce the problem of *generalized zero-shot sign language recognition* (GZSSLR), which allows joint evaluation of the ability of a model to generalize out of the training set for classes with training samples and recognize novel classes.

In order to transfer knowledge across classes, we use textual descriptions of signs taken from a sign language dictionary, together with attribute descriptions gathered from a sign hand shapes dictionary. Using a sign language dictionary for obtaining the class representations has two significant advantages for being (i) readily available, and (ii) prepared by the sign language experts in a detailed way. In this manner, we explore the expressive power of both the textual sign language descriptions and the attributes acquired from dictionaries to form the auxiliary information that is required for recognizing novel sign language classes. Note that, since this auxiliary data is acquired from dictionaries, there is no need for any ad-hoc manual annotations.

To study the ZSSLR problem, we introduce three benchmark datasets. We obtain the first one, ASL-Text , by adapting the ASLLVD dataset [3]. We define the next two ones, MS-ZSSLR-C and MS-ZSSLR-W, by adjusting and manually filtering the MS-ASL [15] dataset. We augment all three datasets with textual and attribute-based class descriptions based on sign language dictionaries.

To tackle the ZSSLR problem, we propose an embedding-based framework that consists of two main components. The first component aims to capture the temporal and spatial patterns, for which we offer alternative architectures based on 3D-CNN, LSTM, and the recently introduced shift-based CNN [16] modules. This component operates over the body and hand regions extracted from

- *E-mail addresses: yunuscanbilge@cs.hacettepe.edu.tr, gcinbis@ceng.metu.edu.tr, nazli@cs.hacettepe.edu.tr*
- *Yunus Can Bilge is with Graduate School of Science and Engineering and Nazli Ikizler-Cinbis is with Computer Engineering Department both from Hacettepe University, Ankara/Turkey. Ramazan Gokberk Cinbis is with the Department of Computer Engineering at METU, Ankara/Turkey.*
- *This work was performed as a part of the Ph.D. studies of the first author, and supported in part by the TUBITAK Grant 116E445.* *© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.*

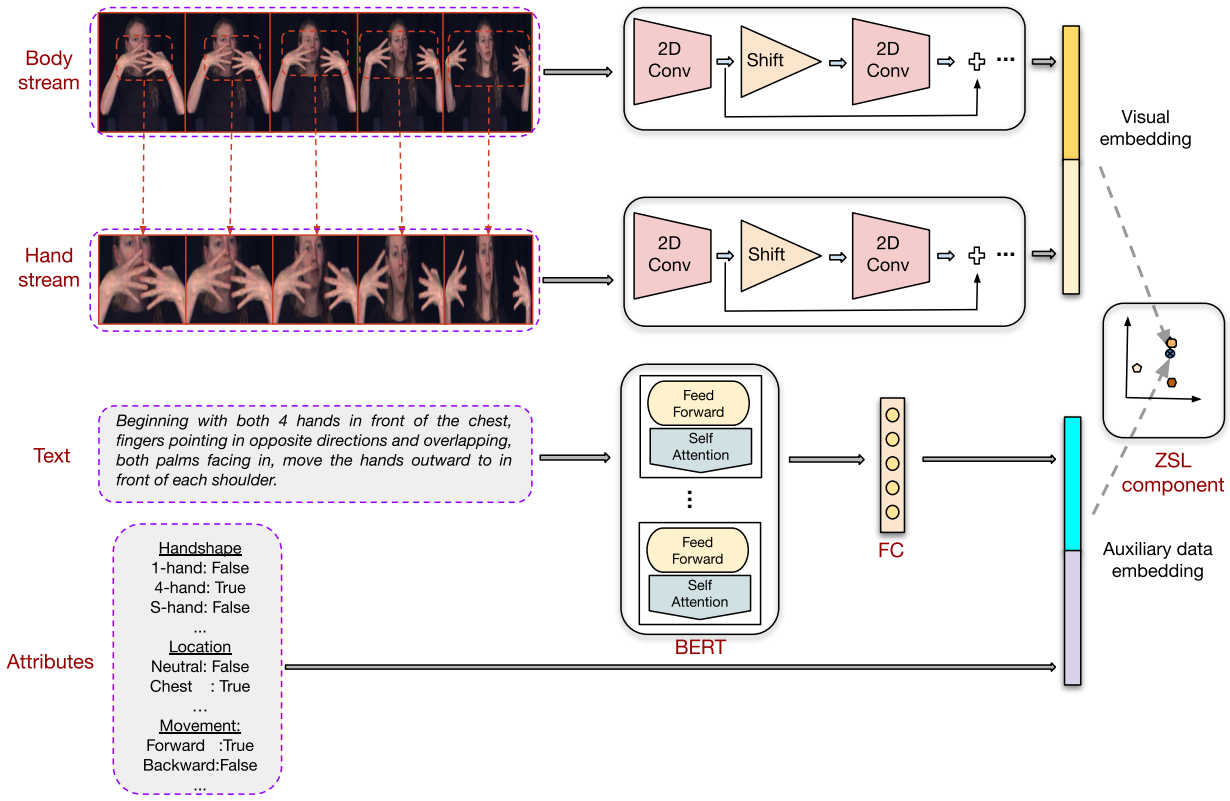

Fig. 1. Illustration of the proposed zero-shot sign language recognition (ZSSLR) approach. Two separate streams are used for both visual and auxiliary class representations. Visual representation is obtained by encoding body and estimated hand streams through spatiotemporal deep architectures. Class embeddings are obtained by encoding textual dictionary definitions and attribute combinations. We learn a compatibility function that links the visual representation to the auxiliary class embeddings for zero-shot recognition.

videos in parallel. The second component, *i.e.*, zero-shot learning (ZSL) component, learns a scoring function for measuring the compatibility between a given visual representation and text and/or attribute-based sign description. We provide an overview in Figure 1. We rigorously evaluate our approach on our three benchmark datasets for both ZSSLR and GZSSLR problems and analyze the results.

A preliminary version of this work has previously appeared in [17], where ZSSLR problem was first formulated and the benchmark dataset ASL-Text was defined. The applied model in [17] was based on 3D-CNNs and LSTMs, the benchmark results were presented on ASL-Text. This work extends it in multiple ways: first, we improve the semantic embeddings of sign classes by collecting new attribute-based representations and show that our attribute-based class definitions yield significant empirical improvements, mainly when used in combination with textual descriptions. Second, we add new temporal shift modules to improve the spatiotemporal video representation. Third, we extend our study by introducing the GZSSLR problem in addition to ZSSLR. Fourth, we introduce two new (G)ZSSLR benchmark datasets. Fifth, we extend our empirical analyses with additional insights into the ZSSLR problems over the three benchmarks, *i.e.* ASL-Text, MS-ZSSLR-W, MS-ZSSLR-C. Finally, we propose two new techniques for analyzing the impact of binary attribute definitions of unseen classes to help us gain insights about attribute-based correct and incorrect zero-shot predictions, respectively.

The remainder of the paper is organized as follows. We start by reviewing the related literature in Section 2. We explain the proposed ZSSLR benchmarks in Section 3. We present our approach to the problem and the attribute analysis technique in Section 4. We present a detailed evaluation and analysis of our models in Section 5. We conclude the paper in Section 6.

2 Related Work

2.1 Sign Language Recognition (SLR)

SLR is a major computer vision task that has been studied for more than three decades [18]. The mainstream SLR approaches can be grouped into two categories: (i) isolated SLR [19], where a single sign is recognized at a time, and (ii) continuous SLR [6], where one or more sentences are recognized over the streaming input. Our work belongs to isolated SLR category as we opt to focus on the study sign recognition with incomplete supervision, isolated from the orthogonal challenges of recognition on streaming data.

Early SLR methods mostly use hand-crafted features in combination with a classifier like support vector machine [18], [20], [21], [22], [23]. Hidden Markov Models (HMM), Conditional Random Fields and neural network based approaches have also been explored to model the temporal patterns [24], [25]. Recently, several deep learning based SLR approaches have been proposed [26], [27], [28].

Despite the relative popularity of the topic, the problem of annotated data scarcity has been seldomly addressed in

TABLE 1
Statistics of the origin datasets; ASLLVD and MS-ASL.

Dataset	Classes	Signer	Videos	Videos per Class	Controlled
RWTH-Boston-50 [22]	50	3	201	4	✓
Purdue ASL [40]	104	14	1834	18	✓
SIGNUM [41]	450	25	12150	27	✓
LSA64 [42]	64	10	3200	50	✓
GSL Corpus [43]	981	2	9810	10	✓
DEVISIGN [44]	2000	8	24000	12	✓
ASLLVD [3]	**3300**	1-6	9800	3	✓
MS-ASL [15]	1000	11-45	**25513**	25	✗

SLR research. Farhadi and Forsyth [29] are the first to study the alignment of sign language video subtitles and signs in action to overcome annotation difficulty. Similarly, [30] proposes to transfer large amounts of labeled avatar data with few labeled human signer videos to spot words in videos. Buehler *et al*. [31] aim to reduce the annotation effort by using the subtitles of British Sign Language TV broadcasts. The idea is to apply Multiple Instance Learning (MIL) to recognize signs out of TV broadcast subtitles. There are also other studies that utilizes the subtitles of TV broadcasts [32], [33]. Pfister *et al*. [33] differ from the two aforementioned MIL studies as these studies track the co-occurrences of lip and hand movements to reduce the search space for visual and textual content mapping. Nayak *et al*. [34] propose to locate signs in continuous sign language sentences using iterated conditional modes. Pfister *et al*. [35] define each sign class with one strongly supervised example and train an SVM based detector out of these one-shot examples. The resulting detector is then used to acquire more training samples from another weakly-labeled data. Koller *et al*. [9] propose a combined CNN and HMM approach to train a model with large but noisy data. Koller *et al*. [36], [37] later suggest using CNN-HMM with LSTM for weakly-labeled noisy data. Albanie *et al*. [38] underline the difficulty in collecting sign language data and propose to extract annotations using keyword spotting in interpreted TV broadcasts. The key idea behind this approach is that signers mostly mouth the words they are signing. Recently, Momeni *et al*. [39] propose a unified Multiple Instance Learning sign spotting framework in continuous sign language videos and evaluate the model in low-shot settings. None of the aforementioned models approach the problem of annotated data scarcity from a zero-shot learning perspective, which can potentially play a central role in modeling larger vocabularies and a large number of languages.

2.2 Sign language datasets

There are many SL datasets available for both isolated [3], [15], [22], [40], [41], [42], [43], [44] and continuous sign language recognition [38], [41], [45], [46], [47]. Table 1 summarizes the details of the currently available public datasets that include isolated sign data to the best of our knowledge. ASLLVD [3] includes the most class count while MS-ASL [15] includes the most in-class sample count. As opposed to most of the available isolated sign language datasets that their recordings are performed in a controlled laboratory environment, MS-ASL [15] covers unconstrained sign recordings concerning large variation in signer, background and positioning. We have chosen ASLLVD [3] and MS-ASL [15] a basis for constructing our zero-shot oriented datasets because of their large class count with challenging real-life recording conditions.

2.3 Zero-Shot Learning (ZSL)

ZSL problems have received significant attention over the past several years, following the pioneering works of Lampert *et al*. [48] and Farhadi *et al*. [49]. The main goal in ZSL is learning to generalize a recognition model for identifying unseen classes. Most of the ZSL approaches rely on transferring semantic information from seen to unseen classes utilizing auxiliary information. Among the possible sources of auxiliary information, attributes and textual representations stand out as powerful and practical options, for which we present a brief overview in the following.

Attribute-based ZSL. Attribute-based ZSL approaches in computer vision rely on class representations based on the information about the relevance of each visual attribute for each class. Attribute information can take many forms, including binary [50], continuous [51] or relative [52]. Attributes have been used for image classification [48], [53], image description [49], [50], object detection [54], [55], and caption generation [56], [57]. We use attributes for high level representation of the sign classes.

Earlier zero-shot recognition approaches use attribute prediction [48], [49], where each attribute is learned independently. More recent works typically focus on directly modeling relations between image features and class attributes, see, *e.g*. [53]. Jayaraman *et al*. [51] propose to model attribute relationships. In recent studies [58], [59], deep learning based approaches to learn visual attributes is proposed.

Text-based ZSL. In text-based ZSL, the semantic space and the representation of a class are constructed out of textual descriptions. Semantic word/sentence vectors and concept ontologies have been studied in this context [60], [61], [62], [63]. Label embedding models are explored to make a connection between seen and unseen classes via semantic representations [53], [64], [65], [66]. Textual descriptions are acquired from various resources; such as Wikipedia [67]. In our case, we use a sign language dictionary as the textual information source. In order to utilize these textual resources, various unsupervised embedding models exist, such as word2vec [68], GloVe [69], and BERT [70].

Zero-shot action recognition. In relation to SLR, ZSL of human actions has been studied. Liu *et al*. [71] are the first to propose attribute based model for recognizing novel actions. Jain *et al*. [72] propose a semantic embedding based approach using commonly available textual descriptions, images, and object names. Xu *et al*. [73] propose a regression based method to embed actions and labels to a common embedding space. Xu *et al*. [74] also use word-vectors as a semantic action embedding space in a transductive setting. Wang *et al*. [75] exploit human actions via related textual descriptions and still images, where the idea is to improve word vector semantic representations with additive information. Habibian *et al*. [76] also propose to learn semantic representations of videos with freely available video and relevant descriptions. Qin *et al*. [66] use error-correcting output codes to overcome the disadvantages of attributes and/or semantic word embeddings of actions.

Hahn *et al*. [77] explore the idea of learning a cross-modal embedding space for textual features of class labels and spatiotemporal features of action videos.

Compared to action recognition, in SLR, even a subtle change in motion and/or handshape can change the entire meaning, which makes SLR a different problem that requires specialized methods for zero-shot recognition. Relatively speaking, gesture recognition is a problem with more similarities to SLR. While there are a couple of recent methods on zero-shot gesture recognition, these methods are limited to either gestures for human-robot interaction [78], or, few classes, *e.g*., 10 hand gesture classes [79]. In this work, we introduce and work on much larger-scale and comprehensive datasets for sign language recognition and incorporate both textual and attribute-based sign representations to our study.

Understanding the influence of class attributes. Recently, there has been a great interest in developing techniques for producing explanations of predictions made by deep architectures. Most main-stream works introduce methods that can be used to estimate the importance of input features, such as LIME [80], DeepSHAP [81], [82], and Grad-CAM [83]. Only very few recent works introduce explanation approaches for zero-shot learning models: Selvaraju *et al*. [84] define a method towards associating neurons with attributes and uses the associations learnt by the model to determine attributes with largest influence on the resulting class scores. Liu and Tuytelaars [85] propose a ZSL model that can produce visual and textual explanations for recognition results. In our work, we are interested in understanding the role of binary class attributes on zero-shot predictions. We note that this differs from analyzing input feature dimensions as attributes can be considered as auxiliary model parameters that are used to infer the class predictors, as opposed to being instance-specific input data.

3 Proposed Datasets

Since there is no available dataset for ZSSLR, we repurpose two existing supervised sign language datasets to create three ZSSLR benchmarks: *ASL-Text*, *MS-ZSSLR-W* and *MS-ZSSLR-C*. These datasets jointly provide a rich experimental setup thanks to significant differences across them, *e.g*., in terms of per-class sample counts and embracing controlled or uncontrolled recording setups. We provide the details of the datasets in the following paragraphs, followed by the explanations about the way we obtain the text-based and attribute-based class definitions.

3.1 ASL-Text dataset

To the best of our knowledge, the ASLLVD dataset [3] is the largest isolated sign language recognition dataset available in vocabulary size. We select the top 250 sign classes, ranked by the number of samples and number of signers per class. We then augment this dataset with the textual definitions of the signs from Webster American Sign Language Dictionary [86] and attribute definitions from American Sign Language Handshape Dictionary [87]. We refer to this new benchmark dataset as *ASL-Text*. Example frames and their textual descriptions from the ASL-Text dataset are presented in the upper half of Figure 2.

In this dataset, the average sequence length for a video sample is 33 frames. While splitting the dataset into three disjoint train, validation and test sets, the classes with the most signer variations, and in-class samples are assigned to the training set. The remaining classes with relatively fewer visual examples are allocated into the validation and test sets to obtain a more realistic experimental setup. The average length of a textual description is 30 words, where the total vocabulary includes 154 distinct words. With only 6.3 samples per class, the dataset provides considerably few training samples, in comparison to commonly used zero-shot image classification datasets, such as aPY [49] and AWA-2 [88], which contain approximately 630 and 750 samples per training class, respectively.

3.2 MS-ZSSLR-W/C datasets

We construct our other ZSSLR benchmarks by adapting the recently introduced *MS-ASL* large-scale dataset [15]. The MS-ASL dataset provides a compelling collection of videos with challenging and realistic sign language sequences.

The dataset is originally collected from public videos with manual captions, descriptions, titles, and subtitles. These videos are processed using automated techniques such as Optical Character Recognition (OCR) to capture the labels of longer video clips, title extraction to be used as the label in the short ones. As a result of these semi-automatic procedures, the dataset contains some artifacts, such as repetitive signs in some single-sign sequences. In addition, most sign classes do not include the standard sign examples but also their sign dialects as well.

We construct two separate benchmarks based on MS-ASL: (i) **MS-ZSSLR-W**(ild), which can be seen as an *in-the-wild* test bed and (ii) **MS-ZSSLR-C**(lean), which provides a more controlled and cleaner benchmark. For both, we first select the largest 200 of the available sign classes in MS-ASL, ranked by the number of samples per class. We then augment the dataset with textual and attribute-based descriptions as we do for constructing the ASL-Text benchmark. We define the resulting dataset, with minimal modifications [1], as the MS-ZSSLR-W benchmark.

The fact that MS-ASL contains language dialect variations, in addition to the within-language variations (*e.g*. photometric, geometric, and person-to-person differences), is considered as an advantage for supervised sign recognition purposes. However, these lingual variations can equally be regarded as *label noise* that reduces the quality of the experimental setup for ZSL purposes: each dialect should typically be accompanied by its own sign class definitions. For instance, there are multiple ways of signing *pizza* in ASL, varying across geographical regions [89]. One popular variation of *pizza* sign mimics shoving a piece of pizza in your mouth, while another highly used variation use *Double-Z* and *A* handshape in conjunction sequentially, which corresponds to completely different signs.[2] Since we have neither the dialect annotations in MS-ASL videos and

1. As an exception, four of the top-200 (*out/outside*, *bus*, *light* and *born*) classes are replaced with top-250 (*money*, *wife*, *week*, *old*) counterparts due to the mismatch between visual sign execution and relevant textual and attribute dictionary descriptions.

2. Variations of the *pizza* sign can be observed in http://www.lifeprint.com/asl101/pages-signs/p/pizza.htm.

TABLE 2
Statistics of the proposed ZSSLR benchmark datasets.

Dataset	Class Count	Train Count	Val. Count	Test Count	Total Count	Avg. per Class	Frame Count	Con-trolled?
ASL-Text	250	1188	151	259	1598	6.3	54,151	✓
MS-ZSSLR-W	200	5862	1153	1846	8861	44.3	922,962	✗
MS-ZSSLR-C	200	7303	1368	1779	10,450	52.2	688,118	✗

nor the per-dialect sign definitions, the ZSL experimental results based on reference sign class definitions mixed with non-compliant sign examples can be misleading.

Based on these observations, we construct a second MS-ASL based benchmark with minimal data and label noise. For this purpose, we have manually verified each video sample and excluded the dialects, and retained only the sign samples that adhere to the reference. We have also split the repetitive sequences into multiple sequences to be used as separate samples. This constitutes our third benchmark, MS-ZSSLR-C. The number of frames per video in these MS-ZSSLR-W/C datasets range from 4 to 298.

A summary of the dataset statistics is provided in Table 2. It can be seen that ASL-Text has the highest class count (250) but the smallest number of samples per class on average (6). MS-ZSSLR-W/C datasets both have 200 classes with 44 and 52 average sample counts per class, respectively, which are approximately 7 times higher than that of ASL-Text. Overall, therefore, MS-ZSSLR-W/C provides a larger number of samples taken in an uncontrolled setting for each class while ASL-Text provides fewer per-class samples taken in a laboratory environment, for a relatively larger number of classes. The dataset splits are explained in Section 5.1.

Figure 2 presents example sign sample frames and the corresponding text and attribute-based descriptions for the ASL-Text and MS-ZSSLR-W/C datasets. It can be observed how textual descriptions and attributes can be complimentary to each other for characterizing signs in terms of handshapes, hand locations, hand movements, palm orientations, combined with some additional details. Our experimental results on ZSSLR and GZSSLR in Section 5 confirm this observation.

3.3 Textual descriptions

We augment all datasets with the textual sign descriptions that are gathered from Webster American Sign Language Dictionary [86]. The textual descriptions of the proposed datasets include the detailed instructions of signs with an emphasis on four basic elements: handshape (*A-hand*, *S-hand*, *5-hand*, *etc.*), the orientation of the palms (*forward, backward, etc.*), movements of the hands (*right, left, etc.*), and the location of hands concerning to the body (*in front of the chest, each side of the body, right shoulder, etc.*). Some descriptions additionally include non-manual cues such as the facial expressions, head movement , and body posture. The average length of a textual description is 30 words with a vocabulary of 154 words and 29 words with a vocabulary of 274 words in the MS-ASL and MS-ZSSLR-W/C datasets, respectively.

Example textual descriptions can be found in Figure 2. Note that hand shapes are described with a specialized vocabulary that involves the terms *F-hand*, *A-hand*, *S-hand*, *5-hand*, *8-hand*, *10-hand*, *open-hand*, *bent-v hand*, *flattened-o hand* [86]. From the example hand shapes shown in Figure 2, it can be seen that the textual sign language descriptions are indeed quite indicative of the ongoing gesture.

3.4 Attribute descriptions

We further annotate the datasets with high-level attributes that are gathered from American Sign Language Hand Shape Dictionary [87]. The attributes highlight four basic features; hand shape (`A-hand`, `S-hand`, `5-hand`, etc.), palm orientation (`in`, `out`, `up`, `down`, `left`, `or right`), hand location (`neutral`, `chest`, `on the shoulder`, `mouth/nose`, `on forehead/eyes`, `ear/temple`), and movement (`up`, `down`, `left`, `right`, `inward`, `outward`, `circular`, `wrist movement`, `finger movement`). We additionally include two relevant attributes: `repetition` indicates whether the hand movement repeats or not, and `one-hand` indicates whether the sign is made by a single hand or both hands. In combination, we obtain a 53-dimensional attribute-based class representations. As can be seen in Figure 2, these selected attributes provide an expressive medium to represent sign characteristics. For instance, `bicycle` sign is described to have `S-hand` handshape with `palms down`, `natural` hand locations, and `circular` movements. Additionally, `circular` movement is repeated and both hands are used in the making of the sign.

4 METHODOLOGY

In this section, we first give a formal definition of the problem and then explain the components of the proposed approach. Finally, we present our binary attribute analysis techniques.

Problem definition. In ZSSLR, there are two sources of information: (i) the *visual domain* $\mathbb{V}$, which consists of sign videos, and, (ii) the *auxiliary information domain* $\mathbb{T}$, which includes the class embeddings based on textual descriptions and/or attributes. At training time, the videos, labels and the sign descriptions are available only for the *seen* classes, $\mathbb{C}_{\text{s}}$. In ZSSLR, at test time, our goal is to correctly classify the examples of novel *unseen* classes, $\mathbb{C}_{\text{u}}$, which are distinct from the seen classes. GZSSLR is essentially the same as ZSSLR except that at test time, the goal expands to correctly classifying the novel examples of both *unseen* classes $\mathbb{C}_{\text{u}}$, and the *seen* classes $\mathbb{C}_{\text{s}}$. In our explanations of the approach below, we focus on the ZSSLR case for simplicity.

The training set $S_{tr} = \{(v_i, c_i)\}_{i=1}^{N}$ consists of N samples where v_i is the i-th training video and $c_i \in \mathbb{C}_{\text{s}}$ is the corresponding sign class label. We assume that we have access to the textual and/or attribute descriptions of each class c, represented by $\tau(c)$. The goal is to learn a zero-shot classifier that can correctly assign each test video to a class in $\mathbb{C}_{\text{u}}$, based on the auxiliary information. We aim to construct a label embedding based zero-shot classification model. For this purpose, we define a *compatibility function* $F(v, c)$ as a mapping from an input video and class pair to a score representing the confidence that the input video v belongs to the class c. Given the compatibility function F, the test-time zero-shot classification function $f : \mathbb{V} \rightarrow \mathbb{C}_{\text{u}}$ is defined as:

$$f(v) = \arg\max_{c \in \mathbb{C}_{\text{u}}} F(v, c). \tag{1}$$

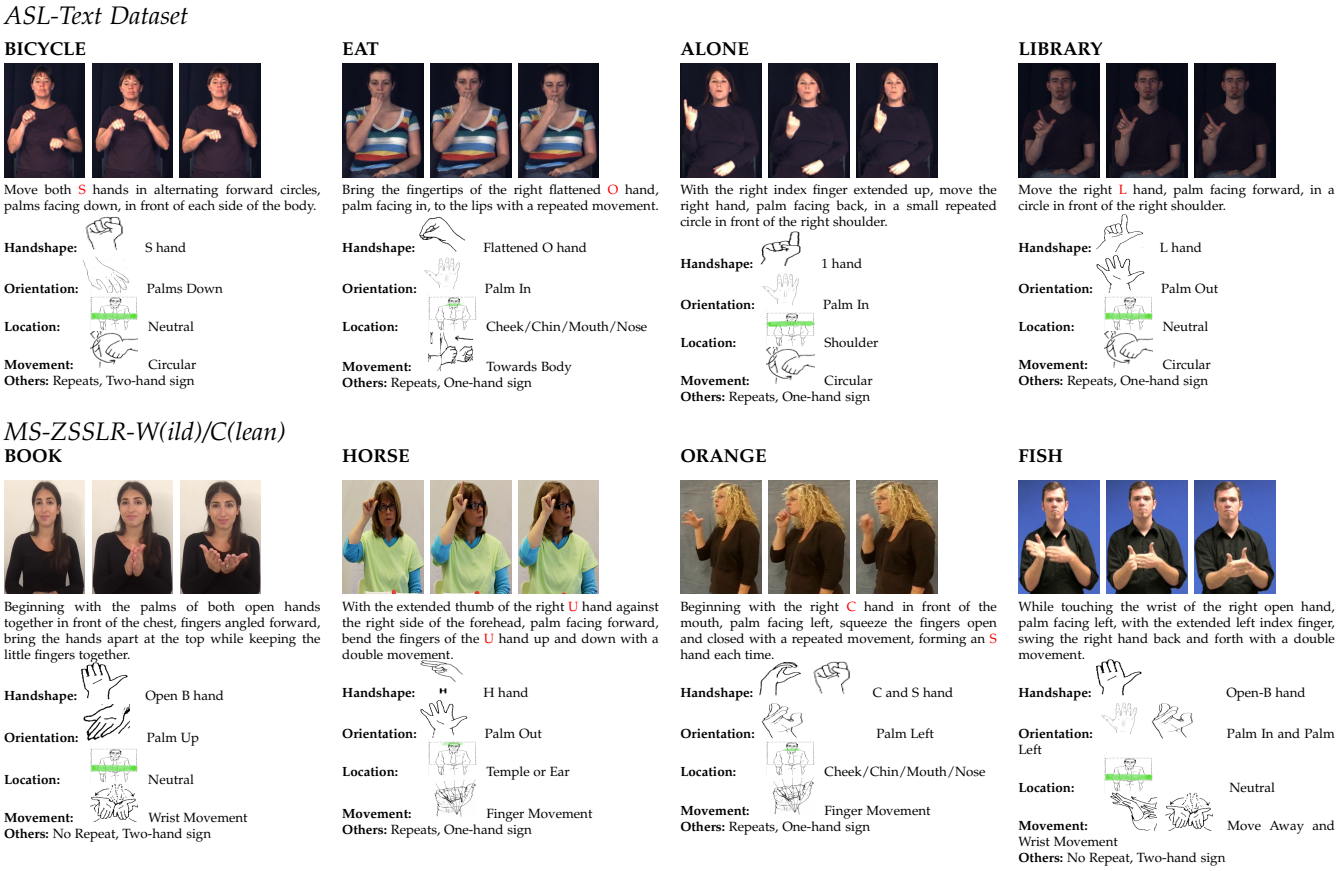

Fig. 2. Example sequences and corresponding textual descriptions from the ASL-Text (upper half) and MS-ZSSLR-W/C (bottom half) datasets. For visualization purposes, only the person regions of the videos are shown.

In this way, we leverage the compatibility function to recognize novel signs at test time. In the case of GZSSLR, the $\arg\max$ operator runs over $\mathbb{C}_u \cup \mathbb{C}_s$ instead.

The performance of the resulting zero-shot sign recognition model directly depends on three factors: (i) video representation, (ii) class representation, and, (iii) the model used as the compatibility function F. An overview of our approach for addressing these issues is shown in Figure 1, and the details are presented in the following sections.

4.1 Spatiotemporal video embedding

In this work, we explore two approaches for obtaining video representations. The first one is by extracting short-term spatiotemporal features using ConvNet features of the video snippets, and then capturing longer-term dynamics through recurrent models. The second approach is capturing spatiotemporal features through the recently proposed temporal shift module (TSM) [16] in an end-to-end manner. We improve the video representation by extracting features in two separate streams: the full frames and hand regions.

Short-term spatiotemporal representation. Short-term spatiotemporal representation is obtained by first splitting each video into 8 frames long snippets and then extracting their features using the pre-trained I3D model [90], a state-of-the-art 3D-ConvNet architecture. I3D model is obtained by adapting a pre-trained Inception model [91] to the video domain and then fine-tuning on the Kinetics dataset. The final video representation is obtained by temporally average pooling the resulting snippet features.

Modeling longer-term dependencies. Average pooling the 3D-CNN features is a well-performing technique for the recognition of non-complex (singleton) actions. Signs, on the contrary, tend to portray more complex patterns that are composed of sequences of multiple basic gestures. In order to capture the transition dynamics and longer-term dependencies across the snippets of a video, we use recurrent models that take an I3D representation sequence as input, and, provide an output embedding. For this purpose, we use a bidirectional LSTM (bi-LSTM) [92] model, and, compare it against the average pooling, LSTM [93] and GRU [94] models.

Temporal shift module. Apart from modeling both short-term and longer-term dependencies in conjunction, we adapt Temporal Shift Module (TSM) [16] to our proposed model as a spatiotemporal video representation module. TSM [16] aims to achieve 3D CNN performance with 2D CNN complexity for large scale video understanding. The module basically shifts information on channels along the temporal dimension in forward and backward to aid information exchange with neighboring frames. For instance, a given 1-D input X with infinite length and a kernel size of 3 with weights of (w_1, w_2, w_3), the shift operation can be defined as:

$$X_i^{-1} = X_{i-1}, X_i^0 = X_i, X_i^{+1} = X_{i+1} \quad (2)$$

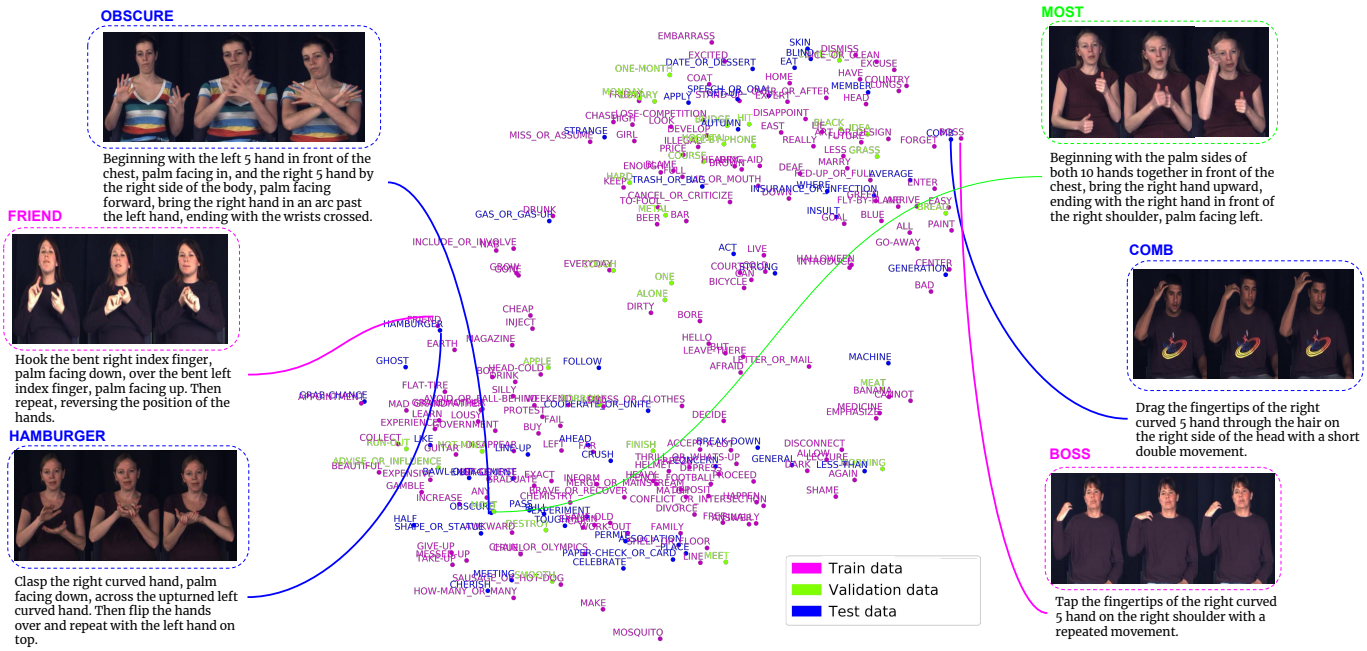

Fig. 3. t-SNE visualization of sign descriptions using BERT- [70] embeddings. Nearby descriptions typically correspond to visually similar signs. Best viewed in color, with zoom.

by shifting the input by -1, 0, +1 in a bidirectional way. In order to get a spatiotemporal video representation, $\phi(v)$, with shifted input, the following multiply-accumulate operation is performed to acquire a visual representation;

$$\phi(v) = w_1 \cdot X^{-1} + w_2 \cdot X^{0} + w_3 \cdot X^{+1} \quad (3)$$

Intuitively, the module first shifts the data and then perform 2D convolution operation. The advantage is that such a shift based approach can be incorporated into any 2D CNN model.

Two-stream video representation. Hands play a central role in expressing signs. In order to encode details of the hand movement, we detect and crop the hand regions using OpenPose [95] and form a hand-only sequence corresponding to each video snippet. As a result, two separate streams are defined, including either TSM [16] or I3D and bi-LSTM networks, over whole video frames and hand regions. The resulting video representations coming from each stream are then concatenated to obtain the final video representation, as illustrated in Figure 1. Note that, both streams are trained together with the compatibility function in an end-to-end fashion.

4.2 Auxiliary knowledge modeling

We use the auxiliary information needed for zero-shot recognition via two main sources; textual descriptions and attributes. The descriptions are both based on sign language dictionaries and include rich information about the signs. We explore the utilization of concatenated textual and attribute descriptions apart from individual usages of them.

Text-based class embeddings. We extract contextualized language embeddings from textual sign descriptions using the state-of-the-art language representation model BERT [70]. BERT architecture basically consists of a stack of encoders; specifically, multi-layer bidirectional transformers [96]. The model's main advantage over the word2vec [68] and glove [69] representations is that BERT model is contextual and the extracted representations of the words change with respect to other words in a sentence. Figure 3 shows the t-SNE visualization of all sign class BERT embeddings. A close inspection to this feature space reveals that classes that appear closer in t-SNE embeddings have indeed similar descriptions. For instance, *friend* and *hamburger* signs are composed of similar motions with different handshapes, *obscure* and *most* signs have similar hand movements but different hand shapes and directions, and, *comb* and *boss* signs include the same repeated movement with different handshape and locations with respect to the body.

Attribute-based class representation. As the attribute based class representations, we use binary attributes such that each class c has fixed-length binary attribute vectors; $\alpha(c) = [a_1^c, a_2^c, ..., a_{53}^c]$ to describe the classes. These attributes are defined based on a handshape dictionary for each sign class that emphasizes four points; handshapes, hand location, palm orientation, hand movement to describe sign making. Figure 2 presents sample attribute based representations. To the best of our knowledge, this is the first study that uses attributes as a common space for sign language recognition.

We also utilize both attributes and textual descriptions in conjunction as the auxiliary information source. To have a more compact representation, we first transform 768 dimensional textual feature vectors into a d_t-dimensional embedding space using a linear layer. We then concatenate the 53 dimensional attribute and d_t dimensional textual feature vectors into $k + 53$ dimensional vectors to represent sign classes. We fix $d_t = 64$ based on the validation set results of ASL-Text dataset. In Section 5.2, we evaluate the corresponding ablation studies using these different information sources and provide details of the d_t tuning experiments.

4.3 Zero-shot learning model

In our work, we adapt a label embedding [53], [97] based formulation to tackle the ZSSLR and GZSSLR problems. More specifically, we use a bi-linear compatibility function that evaluates a given pair of video and class representations:

$$F(v, c) = \phi(v)^T W \rho(c) \tag{4}$$

where $\phi(v)$ is the d-dimensional embedding of the video v, W is the $d \times t$ compatibility matrix, and $\rho(c)$ is the t-dimensional class embedding. The class embedding can be the per-class attribute vector, reduced BERT embedding or their concatenation.

In order to learn this compatibility matrix, we use cross entropy loss with ℓ_2-regularization:

$$\min_W -\frac{1}{N}\sum_{i=1}^{N} \log \frac{\exp\{F(v_i, c_i)\}}{\sum_{c_j \in \mathbb{C}_s} \exp\{F(v_i, c_j)\}} + \lambda \|W\|^2 \tag{5}$$

where λ is the regularization weight. This core formulation is also used in [98], in a completely different ZSL problem. Since the objective function is analogous to the logistic regression classifier, we refer to this approach as *logistic label embedding* (LLE).

In addition to LLE, we also adapt the *embarrassingly simple zero-shot learning* (ESZSL) [65] and *semantic auto-encoder* (SAE) [99] formulations as baselines. The ESZSL formulation utilizes a regression-based loss function in combination with multiple ZSL-specific regularization terms to learn a bi-linear compatibility model between the image and class embeddings. The SAE approach embraces an auto-encoding principle that aims to project image embeddings to the semantic space, and reconstruct back from it. Compared to these two alternatives, LLE takes a more discriminative learning approach simply based on regularized cross-entropy loss over the compatibility scores.

4.4 Analysing binary attribute based class definitions

In our experimental results, presented in the next section, we observe that the auxiliary knowledge based class representations greatly affect the ZSSLR results. To better understand the role of auxiliary knowledge in ZSSLR, binary attributes are particularly valuable for being visually and semantically distinct entities [100]. We note that the case of using dictionary-based binary attributes in ZSSLR differs from the commonly studied case in zero-shot image classification benchmarks, *e.g.* aPY, AWA1, AWA2, CUB and SUN datasets [88], where the class attributes are typically continuous and obtained by averaging per-image manual annotations for experimental purposes. In contrast, we focus on a more realistic setting where the attribute annotations are derived directly from the canonical definitions of signs, without relying on visual sample annotations (of unseen classes). To this end, we aim to evaluate how associating (or not) a class with a binary attribute affects the unseen class predictions. For this purpose, below, we present two techniques for analysing correct and incorrect zero-shot recognition results.

Analysing correct confidence scores. On correct predictions, we aim to understand the relation between the confidence scores and the binary attribute based definitions of classes. Had attributes been continuous entities, we could straightforwardly evaluate the impact of the relation between the k-th attribute and c-th class in terms of partial derivatives:

$$\frac{dp(c|v)}{da_{c,k}} \tag{6}$$

where $a_{c,k}$ is the representation of class c in terms of the k-th attribute and $p(c|v)$ is the posterior probability of class c for the input v. We note that F, used in the definition of $p(c|v)$, depends on $a_{c,k}$ through the use of class embedding function $\rho(c)$. However, since we are interested in binary attribute definitions of sign classes, we cannot directly compute partial derivatives. As a remedy, we define the *flip-difference operator* $\nabla_{c,k}$, analogous to the backward difference based derivative approximations:

$$\nabla_{c,k}[p(v|c)] = p(c|v) - p(c|v; \overline{a_{c,k}}) \tag{7}$$

where we use the notation $p(c|v; \overline{a_{c,k}})$ to express posterior probability of class c for input v when the k-th attribute of the c-th class is flipped (positive to negative, or vice versa) at inference time.

Intuitively, we expect a significant drop in posterior probability when an informative attribute is flipped. Therefore, a relatively large $\nabla_{c,k}[p(v|c)]$ value can be interpreted as the higher importance of the corresponding class attribute relation for some particular input v. To estimate the class-level importance, we take the correctly classified test examples y of some unseen class c, and, average the result of Eq. 7 over the test samples for each attribute separately, resulting in estimated importance scores of attributes in making correct zero-shot predictions.

Analysing zero-shot misclassifications. To better understand incorrect zero-shot predictions, we again use a flip-difference operator based approach to estimate the effect of each attribute. However, in this case instead of simply computing the influence of attributes on posterior probabilities (of wrongly predicted classes), we want to more explicitly answer the question why some other class is being predicted instead of the correct one. For this purpose, we focus on log-derivative of class posterior probabilities given by the ZSL model. More specifically, analogous to partial derivative of log ratios with respect to attributes, we apply the flip-difference operator on the ratios:

$$\nabla_{c,k}[r(c^\star, c^o, v)] = r(c^\star, c^o, v) - r(c^\star, c^o, v; \overline{a_{c^\star,k}}) \tag{8}$$

where $r(c^\star, c^o, v)$ denotes the log-ratio of some (incorrect) class $c^\star$ probability over the correct class c^o for the input v:

$$r(c^\star, c^o, v) = \log \frac{p(c^\star|v)}{p(c^o|v)}. \tag{9}$$

Similarly, $r(c^\star, c^o, v; \overline{a_{c^\star,k}})$ denotes the log ratio when posterior probabilities are computed by flipping the k-th attribute in the definition of the c-th class:

$$r(c^\star, c^o, v; \overline{a_{c^\star,k}}) = \log \frac{p(c^\star|v; \overline{a_{c^\star,k}})}{p(c^o|v; \overline{a_{c^\star,k}})}. \tag{10}$$

Overall, by definition, a large drop in log-ratio value indicates that the probability will shift more quickly from the incorrectly predicted class to the correct one. Therefore, larger influence values in this definition can be interpreted

TABLE 3
Comparison of ZSL models on *ASL-Text*. I3D [90] features are extracted over the whole frames (*i.e*. body stream) only.

Method	visual rep.	Val (30 Classes) top- 1	Test (50 Classes) top-1	top-2	top-5
Random	-	3.3	2.0	4.0	10.0
SAE	body	10.6	8.0	12.0	16.0
ESZSL	body	12.0	**16.9**	**26.0**	**44.4**
LLE	body	**14.1**	11.4	21.2	41.1

TABLE 4
Evaluation of two-stream spatiotemporal representation on ASL-Text. *body* denotes full frame stream, *hand* denotes estimated hand stream. Average pooling is used for aggregating short-term video representations.

Method	visual rep.	Val (30 Classes) top- 1	Test (50 Classes) top-1	top-2	top-5
Random	-	3.3	2.0	4.0	10.0
ESZSL	body	12.0	16.9	26.0	**44.4**
	hand	13.3	11.6	19.6	33.7
	body + hand	14.6	17.1	25.7	43.0
LLE	body	14.1	11.4	21.2	41.1
	hand	15.0	12.6	19.8	37.8
	body + hand	**16.2**	**18.0**	**27.4**	43.8

as the dominant sources of corresponding misclassifications. To obtain confusion-level attribute influence scores, we find the most commonly occurring class confusions, and average the values obtained via Eq. 8 over the test samples of each predicted and ground-truth class tuple.

5 Experiments

In this section, we first describe the implementation details and the experimental setup. [3] Then we present and discuss our experimental results on (G)ZSSLR. Finally, we present our attribute analysis results in Section 5.3.

5.1 Experimental setup and implementation details

We evaluate our models on the proposed zero-shot sign language datasets; ASL-Text, MS-ZSSLR-W, and MS-ZSSLR-C. In our experiments, we fix the number of video frames of each sign video to 32. For LSTMs, we extract 1024-d features from the last average pooling layer of the I3D model using a stride of 4 using every consecutive 8 frames. When modeling the longer temporal context, we set LSTM's or bi-LSTM's initial hidden and cell states to the average pool of each sequence during training.

In this work, we opt to utilize a bidirectional-TSM module with the ResNet-34 backbone [101] pretrained on ImageNet. We utilize TSM architecture by removing its last fully-connected layer and use the output of its preceding the average-pooling layer as the video representation. One of the trade-offs of TSM is that data movement highly increases memory usage, and therefore, we reduce each video sample frame count to 4 frames. We train the TSM model in conjunction with our label embedding approach in an end-to-end manner.

As our text representation, we use the $\mathrm{BERT}_{\mathrm{BASE}}$ model [70] and extract 768-dimensional sentence-based features. Following the description in [70], we concatenate the features from the last four layers of the pretrained Transformer of $\mathrm{BERT}_{\mathrm{BASE}}$ and l_2-normalize them.

We evaluate the models in both the zero-shot learning (ZSL) and generalized zero-shot learning (GZSL) settings. To create dataset splits, we use the largest classes, ranked by the number of in-class samples, for training and the smallest ones for testing. In the ZSL setting, we use 170 classes of ASL-Text and 120 classes of MS-ZSSLR-W/C as the training classes. In all three datasets, the validation and test sets contain 30 and 50 classes, respectively. In the GZSL setting of ASL-Text, train, validation, and test splits contain 170, 200 and 250 classes, respectively. Similarly, the train, validation, and test splits of $MS-ZSSLR-W/C$ datasets contain 120, 150, and 200 classes in the GZSL setting, respectively.

We repeat each experiment 5 times and report the average results. We use top-n accuracy scores normalized by the class sizes. We compute random baseline scores based on 10000 random prediction trials.

3. Datasets and source code is available at https://ycbilge.github.io/towardszslsign.html.

5.2 Experimental results

In the following, we first present our main ZSSLR results, including a through evaluation with ablative studies on the ASL-Text and MS-ZSSLR-C datasets. In Section 5.2.2, we then present additional results on the use of *in-the-wild* benchmark MS-ZSSLR-W(ild) for test-only or train and test purposes, with comparisons to the results on the cleaned version MS-ZSSLR-C(lean). Finally, in Section 5.2.3, we present our generalized zero-shot learning results.

5.2.1 Main results and ablative studies

ZSL formulation. We evaluate three simple and widely used ZSL formulations, namely SAE [99], ESZSL [65], and LLE. In this comparison, for simplicity, we use the average pooled 3D-CNN features of complete frames only as our the video representation and use the ASL-Text dataset. We compare methods in terms of top-1 validation accuracy and top-1, top-2, and top-5 test accuracy scores.

The results are presented in Table 3. We observe that the best performing formulations on the validation and test sets are LLE and ESZSL, respectively. SAE performs worse in all metrics, which can possibly be explained by the fact that SAE incorporates an auto-encoding loss function that aims to reconstruct from video to semantic space with the purpose of reconstructing back from semantic space to video. Such a reconstruction based formulation possibly behaves sub-optimally in the case of high in-class variance and few per-class train samples. Based on these results, we drop SAE from our following experiments.

Two stream representation. In our next experiments, we introduce the second stream based on estimated hand positions and evaluate the effect of using a two-stream representation. For this purpose, we compare three representations obtained by using *body*, *i.e.*, the full-frame input stream, *hand*, or both. The results on the ASL-Text dataset given in Table 4 show that using the two streams in conjunction

TABLE 5
Evaluation of two-stream spatiotemporal representation on MS-ZSSLR-C. *body* denotes full frame stream, *hand* denotes estimated hand stream. Average pooling is used for aggregating short-term video representations.

Method	visual rep.	Val (30 Classes) top - 1	Test (50 Classes) top-1	top-2	top-5
Random	-	3.3	2.0	4.0	10.0
LLE	body	10.8	5.54	9.55	19.5
	hand	10.0	4.05	7.12	15.7
	body + hand	**11.7**	**5.74**	**10.0**	**19.7**

TABLE 6
Comparison of different temporal units with LLE method on ASL-Text and MS-ZSSLR-C datasets. Body and hand streams are used in conjunction, and the textual descriptions are used as auxiliary data.

Temporal	*ASL-Text Dataset* top-1	top-2	top-5	*MS-ZSSLR-C Dataset* top-1	top-2	top-5
AvePool	18.0	27.4	43.8	5.74	10.0	19.7
LSTM	18.2	28	47.2	6.02	10.1	19.7
GRU	19.7	31.8	50.0	6.04	10.5	19.2
bi-LSTM	**20.9**	**32.5**	**51.4**	6.65	10.4	19.9
TSM	16.8	23.2	38.2	**7.01**	**11.5**	**22.1**

improves the top-1 scores on validation and test sets for both ESZSL and LLE formulations. We also observe that LLE outperforms ESZSL using the two-stream representation consistently in all metrics. Due to its greater validation set performance, we fix our ZSL formulation to LLE in the remainder of our experiments.

In Table 4, we also observe that using the body-only representation, while LLE performs better on the validation set, ESZSL outperforms on the test set. This is in contrast to the consistently better performance of LLE over the combined body and hand representation. We believe that this is primarily due to the nature of ZSSLR, where the characteristics of the sign classes can greatly vary between the train, validation, and test sets, which consist of mutually exclusive classes. Despite such performance fluctuations across the ZSL formulations, we, in principle, expect higher correct recognition rates over richer visual representations due to fine-grained distinctions across many sign classes. For example, `friend` and `hamburger` signs are very similar in terms of their overall motion; the hand shape is the only factor that differentiates these two signs. In this respect, richer visual representation based on separate body and hand encoding is expected to be beneficial, which may also contribute to more stable ZSL results.

In Table 5, we present the evaluation of two-steam representation on the MS-ZSSLR-C dataset. The results consistently show that using body and hand stream representations jointly improves the results over body-only and hand-only results, on both validation and test sets. Overall, the results on both datasets confirm that hand and body streams provide complementary information for zero-shot recognition of signs.

Temporal models. We further evaluate the effect of sequential temporal modeling using recurrent architectures and the TSM model [16], and compare to the simple temporal average pooling scheme. As the recurrent model alternatives, we

TABLE 7
Evaluation of auxiliary knowledge sources on ASL-Text and MS-ZSSLR-C datasets, using LLE model with body and hand streams.

Temporal	Auxiliary knowledge	**ASL-Text** top-1	top-2	top-5	**MS-ZSSLR-C** top-1	top-2	top-5
bi-LSTM	Text	20.9	32.5	51.4	6.65	10.4	19.9
	Attr	23.7	38.6	59.2	12.9	22.4	41.8
	Attr + Text	**31.3**	**46.8**	**66.0**	13.5	24.1	**45.2**
TSM	Text	16.8	23.2	38.2	7.01	11.5	22.1
	Attr	21.8	32.2	55.7	13.1	23.3	41.4
	Attr + Text	25.6	37.8	57.8	**14.7**	**24.6**	43.4

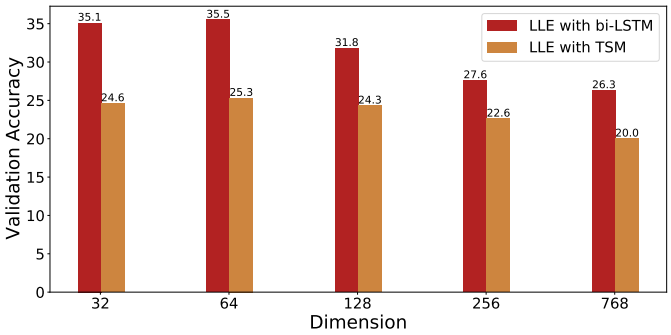

Fig. 4. ASL-Text validation set accuracy as a function of textual embedding dimensionality, in the case of combined text and attribute class embeddings. The 768 bin corresponds using the original textual embeddings, without applying a reduction layer.

experiment with LSTM [93], GRU [94], and bi-LSTM [92] based temporal models. For all cases, we use both hand and body video streams, textual description based class embeddings, and the LLE model.

The results for both ASL-Text and MS-ZSSLR-C are presented in Table 6. Overall, we observe that, compared to temporal average pooling, the framework benefits from the introduction of explicit temporal models. On ASL-Text, the best performing model is bi-LSTM, and TSM-based temporal model does not generalize well. In contrast, on MS-ZSSLR-C, we observe that TSM performs the best in all metrics. These results suggest that TSM performs better in the presence of a larger training set, which is probably related to the fact that the TSM representation is learned in an end-to-end fashion, without utilizing a pre-trained I3D model as in bi-LSTM.

Our overall proposed framework reaches a top-1 normalized accuracy of 20.9 and 7.01, and top-5 normalized accuracy of 51.4 and 19.9 on the test set of the ASL-Text and MS-ZSSLR-C datasets, respectively.

Auxiliary class knowledge. In our final analysis, we explore the options for auxiliary class knowledge sources. We evaluate models using textual descriptions, attributes, or both jointly as class embeddings. We evaluate each option with bi-LSTM and TSM temporal models, following our observations in the preceding experiments.

The empirical comparison of class embedding options is presented in Table 7. We first observe that attribute embeddings perform significantly better than the textual descriptions, for both bi-LSTM and TSM based results on both datasets, in all metrics. We also observe that using attributes and textual descriptions jointly improves attribute-only and

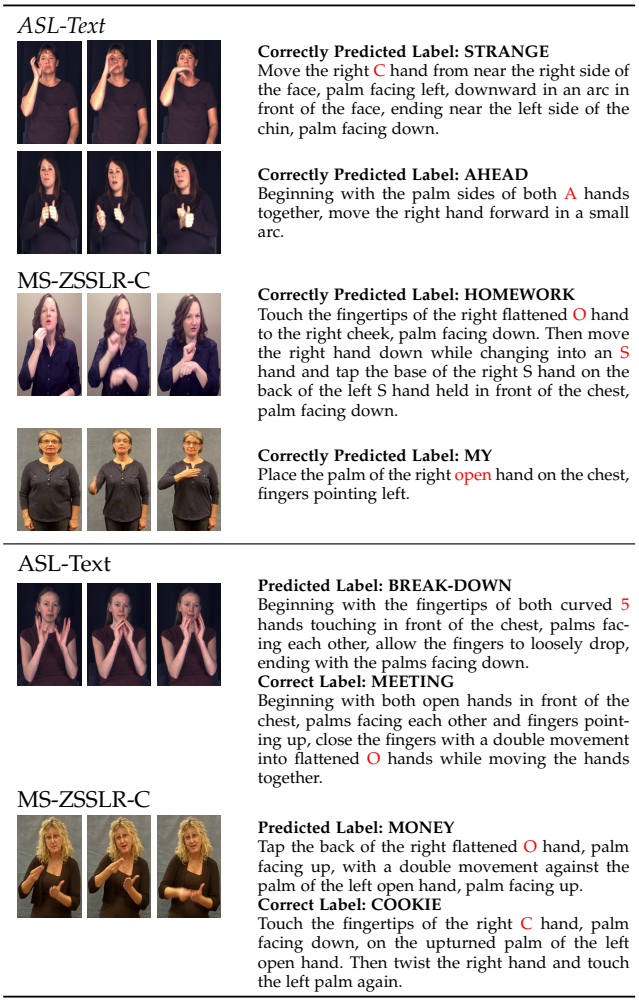

Fig. 5. Correct (first four rows) and incorrect (last two rows) zero-shot prediction examples on the ASL-Text and MS-ZSSLR-C datasets. The textual descriptions of the predicted and ground-truth (separately for incorrect predictions) classes are shown for each case.

text-only results significantly in all cases, reaching a top-1 score of 31.3 on ASL-TEXT using bi-LSTM and 14.7 on MS-ZSSLR-C using TSM. Similarly, using joint embeddings instead of text-only embeddings, the best top-5 scores impressively improves from 51.4 to 66.0 on ASL-Text and 19.9 to 45.2 on MS-ZSSLR-C.

These results indicate that the sign class attributes that we define, carry valuable information and two auxiliary knowledge sources are complementary to each other. The results also show the central importance of class embeddings in zero-shot recognition. Therefore, a promising research direction, in addition to representation and recognition modeling, can be the exploration of new ways for obtaining and leveraging auxiliary knowledge sources.

An important detail in the use combined class embeddings is the trainable linear layer that we use to transform the original 768-dimensional textual embeddings into smaller d_t dimensional vectors, before concatenating with the attribute embeddings. Figure 4 shows the effect of d_t choice on the ASL-Text validation scores, based on which we set $d_t = 64$. In the figure, the 768 bin corresponds using the original textual embeddings directly. The results show that the both LLE with bi-LSTM and LLE with TSM models clearly benefit from the introduction of the transformation layer and both models achieve the highest validation results with $d_t = 64$. With this setting, the accuracy scores of the bi-LSTM and TSM based models improve from 26.3 to 35.5, and from 20.0 to 25.3, respectively.

Qualitative results. Figure 5 presents several correct and incorrect zero-shot prediction examples for the ASL-Text and MS-ZSSLR-C datasets. Through the correct prediction examples, we observe that the model is able to perform well on sign classes with both simple and complex patterns and descriptions. On the incorrect prediction examples, we observe that the descriptions of the confused class pairs are similar to each other, highlighting the difficulty of the tackled problem.

5.2.2 *Exploration of MS-ZSSLR-W(ild)*

As explained in Section 3, we obtain MS-ZSSLR-C(lean) by manually filtering samples from the MS-ZSSLR-W(ild) dataset, according to their adherence to their canonical sign definitions. While the variations of signs deserve an alternative definitions of the signs for a reliable analysis of the ZSSLR models, MS-ZSSLR-W still provides an interesting test bed for understanding the effects of such deviations from the canonical sign definitions. For this reason, we explore MS-ZSSLR-W through three new experimental protocols where we use MS-ZSSLR-W samples for training only, testing only or for both training and testing, and compare against the clean-only setting. For these experiments, we use the configuration that yields the highest results on MS-ZSSLR-C, *i.e.*, LLE zero-shot learning formulation using TSM as the visual representation and attributes+text embeddings as the class embeddings.

The corresponding results are presented in Table 8. We observe that comparing *Clean Train w/ Clean Evaluation* results to *Clean Train w/Wild Evaluation* results, the model trained on clean samples, yields a lower accuracy on wild samples. This is not surprising, since the wild samples includes different dialects that are incompatible with the textual class definitions. Note that the results in Clean vs Wild evaluation setups are not directly comparable, since these tests sets are not identical. When we compare the *Clean Train w/Clean Evaluation* versus *Wild Train w/ Clean Evaluation*, *i.e.* the training set is switched to wild setting over the same test set of clean examples, we observe that incorporating dialects into the training harms the recognition performance and the model trained on MS-ZSSLR-C dataset is more successful. The results over the *Wild Evaluation* follows a similar pattern, and the model trained on the clean samples yields better performance.

These results show the importance of compatibility across visual data and auxiliary representations for both training and testing purposes. The samples incompatible with the class definitions not only reduce the recognition success rates when they appear during the test time, but they also degrade the model quality when used as training samples. As pointed out by these results, modeling sign languages with dialects, especially in zero-shot setting, is an important open problem. We believe that the MS-ZSSLR-C/W dataset pair provides a valuable benchmark for future research in this direction.

TABLE 8
Evaluation over the MS-ZSSLR-C and MS-ZSSLR-W dataset combinations.

MS-ZSSLR-C Clean Train w/ Clean Evaluation			MS-ZSSLR-C ∩ MS-ZSSLR-W Clean Train w/ Wild Evaluation			MS-ZSSLR-W Wild Train w/ Wild Evaluation			MS-ZSSLR-W ∩ MS-ZSSLR-C Wild Train w/ Clean Evaluation		
top-1	top-2	top-5	top-1	top-2	top-5	top-1	top-2	top-5	top-1	top-2	top-5
14.7	**24.6**	**43.4**	10.3	17.8	32.6	8.7	15.7	30.3	10.3	18.6	36.0

TABLE 9
Generalized zero-shot learning (GZSL) results on ASL-Text.

Method	Auxiliary	Val (200 Classes)	Test (250 Classes)		
		top-1	top-1	top-2	top-5
Random	-	0.66	0.5	1.0	2.5
f-clswgan [102]	Text	18.4	11.6	17.9	34.1
	Attribute	27.8	21.5	29.4	44.9
	Attribute + Text	28.2	24.8	37.0	50.5
tfvaegan [103]	Text	22.0	14.8	19.3	31.6
	Attribute	25.4	24.1	31.8	46.4
	Attribute + Text	26.8	26.2	34.9	53.0
LLE with bi-LSTM	Text	29.7	22.5	32.5	45.6
	Attribute	30.5	23.5	33.8	48.4
	Attribute + Text	**38.4**	**26.9**	**39.5**	**55.8**

TABLE 10
Generalized zero-shot learning (GZSL) results on MS-ZSSLR-C.

Method	Auxiliary	Val (150 Classes)	Test (200 Classes)		
		top-1	top-1	top-2	top-5
Random	-	0.66	0.5	1.0	2.5
f-clswgan [102]	Text	28.9	21.5	30.4	42.0
	Attribute	32.0	29.8	38.0	49.3
	Attribute + Text	33.4	31.0	39.5	51.6
tfvaegan [103]	Text	27.0	22.7	28.3	38.6
	Attribute	33.5	32.6	40.4	51.7
	Attribute + Text	34.5	34.2	41.9	52.2
LLE with TSM	Text	45.8	33.4	40.5	48.8
	Attribute	45.7	33.9	42.2	52.9
	Attribute + Text	**46.3**	**34.7**	**42.6**	**53.4**

5.2.3 Generalized-ZSSLR results

Finally, we present the results for generalized zero-shot learning on the ASL-Text and MS-ZSSLR-C datasets. On both datasets, we use the best-performing spatiotemporal representation according to the ZSL results, *i.e.*, bi-LSTM for ASL-Text and TSM for MS-ZSSLR-C, with LLE formulation. On both datasets, we evaluate attribute-only, text-only and attribute-text combination based class embeddings. In addition, we also adapt and evaluate two recent generative GZSL approaches, f-clswgan [102] and tfvaegan [103], as generative approaches tend to yield more comparable seen and unseen class scores for GZSL. The f-clswgan [102] approach combines a Wasserstein GAN [104] with a classification loss to learn a class embedding conditional feature generating model. The tfvaegan [103] approach involves learning a variational auto-encoder [105] with adversarial training and a feedback mechanism, and similarly learns a conditional feature generation model. To adapt both approaches, we use our pre-trained, best-performing models to extract spatio-temporal representations, and train the conditional generative models using these features. To obtain the final classifier, we train supervised models over real and generated examples following the implementation details of [102] and [103].

The GZSL results in terms of top-k accuracy metrics for ASL-Text and MS-ZSSLR-C are presented in Table 9 and Table 10, respectively. On ASL-Text dataset, when using only attribute based class embeddings, tfvaegan [103] yields the highest accuracy scores. With text-only, and combined embeddings, the LLE model outperforms the generative approaches. Using the LLE formulation, we observe that text and attribute-based class representations yield comparable performance on the validation and test sets of both datasets, with relatively better results for the attribute-based class embeddings. We observe significant performance gains, especially on the ASL-Text dataset, when two embeddings are used in combination in all formulations.

In the GZSL setting, the separate accuracy values on the *seen* and *unseen* class samples are of interest. For this reason, we present the seen and unseen top-k class accuracy scores, and their harmonic means (adapted from [88]) on ASL-Text and MS-ZSSLR-C test sets in Table 11. In these results, we observe that unseen class accuracy values are much lower that the seen ones, similar to the case in zero-shot image classification benchmarks [88]. In contrast to the image classification benchmarks, however, we do not see a clear advantage of the generative approaches in terms of unseen class recognition performances, with mixed results across the datasets. These results suggest that the ZSSLR problem has unique challenges, requiring dedicated efforts for developing representation and recognition models.

5.3 Attribute usage analysis

Our experimental results show that the auxiliary class representation has a major influence on the success of zero-shot sign language recognition models. In this section, we present the results of our attribute analysis techniques defined in Section 4.4. Throughout this section, we use the test set predictions on the MS-ZSSLR-C dataset using LLE model based on TSM features and attribute-only class representations. We present our observations for correct and incorrect predictions in the following paragraphs.

Observations on correct zero-shot predictions. We present binary attribute confidence scores for five randomly chosen unseen classes in Figure 6. Here, we take the correctly classified samples of these classes and compute the average influence of each attribute according to Eq. 7. Each row shows the influence scores of the attributes for one class. By definition, the influence scores can be interpreted as the degree of association between the attribute and class pairs.

From the results, we observe that our model strongly associates certain attributes with the relevant classes. For example, *chest* attribute appears to have strong positive effects on the *my* and *love* class predictions, and, negative effects on the other classes, all of which are consistent with the auxiliary knowledge on class-attribute relations. Similarly,

TABLE 11
Generalized zero-shot learning accuracy statistics on ASL-Text and MS-ZSSLR-C test sets.

Method	Dataset	Overall Acc.	Seen			Unseen			Harmonic		
		top-1	top-1	top-2	top-5	top-1	top-2	top-5	top-1	top-2	top-5
f-clswgan [102]	ASL-Text	24.8	33.3	48.2	64.6	6.7	13.2	20.7	11.1	20.7	31.3
tfvaegan [103]		26.2	35.2	45.8	66.8	7.1	11.5	23.5	11.8	18.3	34.7
LLE with bi-LSTM		26.9	37.0	53.3	72.4	5.5	10.4	20.3	9.5	17.4	31.7
f-clswgan [102]	MS-ZSSLR-C	31.0	49.2	60.9	75.1	3.8	7.4	16.2	7.0	13.1	26.6
tfvaegan [103]		34.2	54.2	63.8	74.4	4.3	8.9	18.8	7.9	15.6	30.0
LLE with TSM		34.7	54.6	64.6	76.2	4.8	9.7	19.1	8.8	16.8	30.5

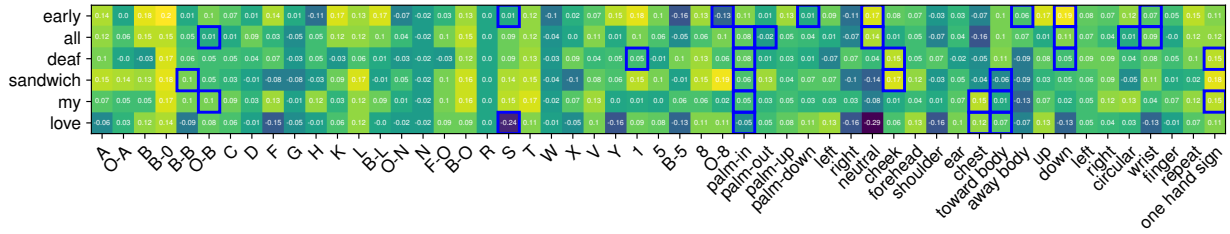

Fig. 6. Influence of the attributes on the confidence scores of correct classified samples of five randomly chosen classes. Each row corresponds to a class and each column corresponds to an attribute. Shown numerical values indicate the average influence of an attribute on the corresponding class confidence scores. Thick boxes show the positive attribute relations according to the attribute based class definitions. Lighter colors correspond to higher values. Best viewed in color with zoom.

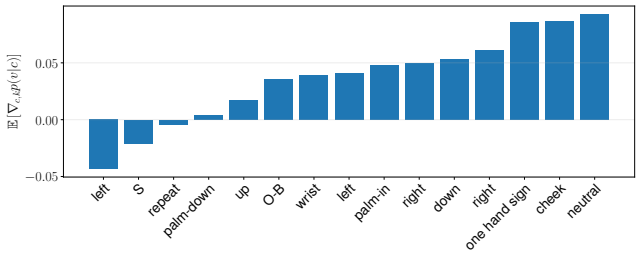

Fig. 7. Influence of positive attributes on the confidence scores of correctly classified test samples, averaged over classes. Only the attributes positively associated with at least 10 unseen classes are shown.

cheek and *neutral* attributes influence scores show consistent patterns with the class-attribute relation definitions. The influence scores also point to certain cases where the model seems to make incorrect, or unexpected, associations between the visual patterns and class attributes. For example, *O-8* (stands for *Open-8*) and *8* appears to be negatively correlated with the attribute definitions of the *early* class. We believe that this is due to subtle differences, where the model is prone to making incorrect generalizations.

To further understand the overall impact of the attributes, for each attribute, we choose the unseen classes that are positively affiliated with that attribute and average the attribute influence values over these classes. In Figure 7, we show the results for the subset of attributes where there exists at least 10 such positively affiliated class according to the class attribute definitions. Consistent with our aforementioned observations, we observe that *neutral* appears to be a strong positive contributor to the related correct class confidence scores, similar to many other attributes such as *cheek* (stands for *cheek/chin/mouth/nose*) and *one hand sign*. While we would ideally expect to see only such positive values for all attributes, few attributes, namely *left* and *S* (and marginally *repeat*), negatively affect the predictions on average. One of the factors in poor modeling of these attributes is possibly their rarity among the training examples, as *left* and *S* are affiliated respectively with only 8 and 10 classes, in contrast to 66 affiliations of *neutral*. Overall, these analysis results show that the ZSSLR model tends to learn meaningful relations, but there is room for improvement in building more explainable and robust models.

Observations on zero-shot misclassifications. Finally, we look into the influence of class-attribute definitions on zero-shot misclassifications. For this analysis, we find four most commonly confused ground-truth - predicted class tuples, and compute the attribute influence scores based on log-ratio of class posteriors, according to Eq. 8. We present the results in Figure 8, where each row corresponds to a particular class confusion case. Thick boxes and circles indicate that the corresponding attribute belongs to the corresponding predicted and ground-truth classes, respectively. We note that the influence scores here are based on the rate of probability transition from incorrectly predicted classes to the correct ones. Therefore, a larger value indicates a relatively bigger role of an attribute on the resulting misclassification.

In Figure 8, we observe large influence scores in some attributes that are positively defined for both the predicted and the ground-truth class, as expected. This is particularly the case with the *neutral* attribute, which seems to be modelled accurately according to our preceding observations, as well. However, this is not always the case, which suggests that the model suffers from poor (and unclear) attribute recognitions in these problematic misclassification cases. It is also noticeable that *B-0*, which stands for *Baby 0*, is inactive for all classes, yet activating this attribute in the definition of the assigned classes would yield great drops. It turns out that this attribute is positively defined for only one training class, explaining why the model might be behaving instable as a function of this attribute. Overall, these results suggest that building ZSL models with more robust (implicit) attribute predictors can potentially be a fruitful future research direction.

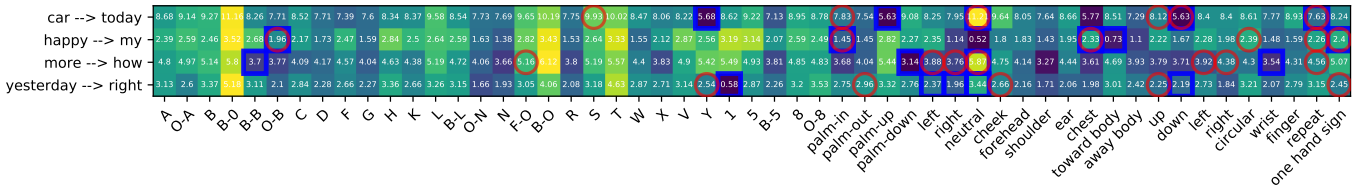

Fig. 8. Influence of the attributes on the confidence scores of zero-shot misclassifications. Each row corresponds to a misclassification case, denoted as *ground-truth class* → *predicted class*. Shown numerical values indicate the average influence score according to Eq. 8. Thick boxes and circles indicate that the corresponding attribute belongs to the corresponding predicted and ground-truth classes, respectively. Lighter colors indicate to higher values, colors in each row are predicted independently for better visualization. Best viewed in color with zoom.

6 Conclusion

This paper explores the problem of ZSSLR and GZSSLR. We present three benchmark datasets for this novel problem by augmenting two large ASL dataset with sign language dictionary descriptions and attributes. Our proposed framework builds upon the idea of using these auxiliary texts and attributes as additional sources of information to recognize unseen signs. We propose an end-to-end trainable ZSSLR method that focuses hand and full body regions via several temporal modeling approaches and learns a compatibility function via label embedding. Overall, the experiments yield promising results on zero-shot recognition of signs. Nevertheless, the acquired accuracy levels are relatively low compared to other ZSL domains, pinpointing a substantial need for further exploration of this direction. The observations based on our binary attribute analysis methodologies suggest that one of the fundamental challenges in ZSSLR is to build models that accurately relate visual patterns to the elements of class definitions, instead of relying on spurious correlations, in making zero-shot predictions.

In our work, we observe two fundamental challenges inherent to the (G)ZSSLR task. First, some of the differences across sign descriptions are subtle, both visually and textually. Second, many dialects exists; even the same sign can be expressed in many different forms. Future work should try to address these issues, possibly by developing visual representations and recognition models that can better capture fine details. Incorporation of additional visual cues, *e.g.* facial expressions and pose information, are also likely to be beneficial. We believe that (G)ZSSLR is an important research direction towards building large-vocabulary sign language recognition systems.

References

[1] Y. Wu and T. S. Huang, "Vision-based gesture recognition: A review," in *International Gesture Workshop*, 1999, pp. 103–115.

[2] W. C. Stokoe Jr, "Sign language structure: An outline of the visual communication systems of the american deaf," *Journal of deaf studies and deaf education*, vol. 10, no. 1, pp. 3–37, 2005.

[3] C. Neidle, A. Thangali, and S. Sclaroff, "Challenges in development of the american sign language lexicon video dataset (asllvd) corpus," in *Proc. 5th Workshop on the Representation and Processing of Sign Languages: Interactions between Corpus and Lexicon, Language Resources and Evaluation Conference (LREC) 2012*, 2012.

[4] C. Valli and C. Lucas, *Linguistics of American sign language: An introduction*. Gallaudet University Press, 2000.

[5] C. Lucas and R. Bayley, "Variation in sign languages: Recent research on asl and beyond," *Language and Linguistics Compass*, vol. 5, no. 9, pp. 677–690, 2011.

[6] N. Cihan Camgoz, S. Hadfield, O. Koller, H. Ney, and R. Bowden, "Neural sign language translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 7784–7793.

[7] N. C. Camgoz, S. Hadfield, O. Koller, and R. Bowden, "Subunets: End-to-end hand shape and continuous sign language recognition," in *Proc. IEEE Int. Conf. on Computer Vision*, 2017, pp. 3075–3084.

[8] O. Koller, J. Forster, and H. Ney, "Continuous sign language recognition: Towards large vocabulary statistical recognition systems handling multiple signers," *Comput. Vis. Image Understand.*, vol. 141, pp. 108–125, 2015.

[9] O. Koller, O. Zargaran, H. Ney, and R. Bowden, "Deep sign: hybrid cnn-hmm for continuous sign language recognition," in *British Machine Vision Conference*, 2016.

[10] S. Stoll, N. C. Camgoz, S. Hadfield, and R. Bowden, "Sign language production using neural machine translation and generative adversarial networks," in *British Machine Vision Conference*, 2018.

[11] M. P. Lewis and F. Gary, "Simons, and charles d. fennig (eds.). 2013," *Ethnologue: Languages of the world,*, pp. 233–62, 2015.

[12] C. Valli, *The Gallaudet Dictionary of American Sign Language*. Gallaudet University Press, 2006.

[13] D. Brien, *Dictionary of British Sign Language: English*. Faber & Faber, 1992.

[14] T. A. Johnston, *Signs of Australia: A new dictionary of Auslan (Australian Sign Language)*. North Rocks Press, 1998.

[15] H. Vaezi Joze and O. Koller, "MS-ASL: A large-scale data set and benchmark for understanding american sign language," in *The British Machine Vision Conference (BMVC)*, September 2019.

[16] J. Lin, C. Gan, and S. Han, "Tsm: Temporal shift module for efficient video understanding," in *Proc. IEEE International Conference on Computer Vision*, 2019, pp. 7083–7093.

[17] Y. C. Bilge, N. Ikizler-Cinbis, and R. G. Cinbis, "Zero-shot sign language recognition: Can textual data uncover sign languages?" in *Proc. British Machine Vision Conference (BMVC)*, 2019.

[18] S. Tamura and S. Kawasaki, "Recognition of sign language motion images," *Pattern recognition*, vol. 21, no. 4, pp. 343–353, 1988.

[19] H. Wang, X. Chai, X. Hong, G. Zhao, and X. Chen, "Isolated sign language recognition with grassmann covariance matrices," *ACM Transactions on Accessible Computing (TACCESS)*, vol. 8, no. 4, p. 14, 2016.

[20] M. B. Waldron and S. Kim, "Isolated asl sign recognition system for deaf persons," *IEEE Transactions on rehabilitation engineering*, vol. 3, no. 3, pp. 261–271, 1995.

[21] M. W. Kadous *et al.*, "Machine recognition of auslan signs using powergloves: Towards large-lexicon recognition of sign language," in *Proc. Workshop on the Integration of Gesture in Language and Speech*, vol. 165, 1996.

[22] M. Zahedi, D. Keysers, T. Deselaers, and H. Ney, "Combination of tangent distance and an image distortion model for appearance-based sign language recognition," in *Joint Pattern Recognition Symposium*, 2005, pp. 401–408.

[23] H. Cooper and R. Bowden, "Sign language recognition using boosted volumetric features," in *Proc. IAPR Conference on Machine Vision Applications*, 2007, pp. 359–362.

[24] K. Grobel and M. Assan, "Isolated sign language recognition using hidden markov models," in *IEEE International Conference on Systems, Man, and Cybernetics. Computational Cybernetics and Simulation*, vol. 1, 1997, pp. 162–167.

[25] C.-L. Huang and W.-Y. Huang, "Sign language recognition using

model-based tracking and a 3d hopfield neural network," *Machine vision and applications*, vol. 10, no. 5-6, pp. 292–307, 1998.

[26] N. C. Camgoz, O. Koller, S. Hadfield, and R. Bowden, "Sign language transformers: Joint end-to-end sign language recognition and translation," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2020, pp. 10 023–10 033.

[27] D. Li, C. Rodriguez, X. Yu, and H. Li, "Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison," in *The IEEE Winter Conference on Applications of Computer Vision*, 2020, pp. 1459–1469.

[28] B. Saunders, N. C. Camgöz, and R. Bowden, "Progressive transformers for end-to-end sign language production," in *European Conference on Computer Vision (ECCV)*, 2020.

[29] A. Farhadi and D. Forsyth, "Aligning asl for statistical translation using a discriminative word model," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, vol. 2, 2006, pp. 1471–1476.

[30] A. Farhadi, D. Forsyth, and R. White, "Transfer learning in sign language," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2007, pp. 1–8.

[31] P. Buehler, A. Zisserman, and M. Everingham, "Learning sign language by watching tv (using weakly aligned subtitles)," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 2961–2968.

[32] D. Kelly, J. Mc Donald, and C. Markham, "Weakly supervised training of a sign language recognition system using multiple instance learning density matrices," *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, vol. 41, no. 2, pp. 526–541, 2011.

[33] T. Pfister, J. Charles, and A. Zisserman, "Large-scale learning of sign language by watching tv (using co-occurrences)." in *British Machine Vision Conference*, 2013.

[34] S. Nayak, S. Sarkar, and B. Loeding, "Automated extraction of signs from continuous sign language sentences using iterated conditional modes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 2583–2590.

[35] T. Pfister, J. Charles, and A. Zisserman, "Domain-adaptive discriminative one-shot learning of gestures," in *Proc. European Conf. on Computer Vision*, 2014, pp. 814–829.

[36] O. Koller, S. Zargaran, and H. Ney, "Re-sign: Re-aligned end-to-end sequence modelling with deep recurrent cnn-hmms," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4297–4305.

[37] O. Koller, C. Camgoz, H. Ney, and R. Bowden, "Weakly supervised learning with multi-stream cnn-lstm-hmms to discover sequential parallelism in sign language videos," *IEEE transactions on pattern analysis and machine intelligence*, 2019.

[38] S. Albanie, G. Varol, L. Momeni, T. Afouras, J. S. Chung, N. Fox, and A. Zisserman, "BSL-1K: Scaling up co-articulated sign language recognition using mouthing cues," in *ECCV*, 2020.

[39] L. Momeni, G. Varol, S. Albanie, T. Afouras, and A. Zisserman, "Watch, read and lookup: learning to spot signs from multiple supervisors," in *Proc. Asian Conference on Computer Vision*, 2020.

[40] R. Wilbur and A. Kak, "Purdue rvl-slll american sign language database. school of electrical and computer engineering technical report," TR-06-12, Purdue University, W. Lafayette, IN 47906, Tech. Rep., 2006.

[41] U. Von Agris, M. Knorr, and K.-F. Kraiss, "The significance of facial features for automatic sign language recognition," in *IEEE International Conference on Automatic Face & Gesture Recognition*, 2008, pp. 1–6.

[42] F. Ronchetti, F. Quiroga, C. A. Estrebou, L. C. Lanzarini, and A. Rosete, "Lsa64: an argentinian sign language dataset," in *Congreso Argentino de Ciencias de la Computación (CACIC).*, 2016.

[43] E. Efthimiou and S.-E. Fotinea, "GSLC: creation and annotation of a greek sign language corpus for hci," in *Int. Conf. on Universal Access in Human-Computer Interaction*, 2007.

[44] X. Chai, H. Wang, and X. Chen, "The devisign large vocabulary of chinese sign language database and baseline evaluations," *Technical report VIPL-TR-14-SLR-001. Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS*, 2014.

[45] A. Roussos, S. Theodorakis, V. Pitsikalis, and P. Maragos, "Dynamic affine-invariant shape-appearance handshape features and classification in sign language videos," *The Journal of Machine Learning Research*, vol. 14, no. 1, pp. 1627–1663, 2013.

[46] J. Forster, C. Schmidt, O. Koller, M. Bellgardt, and H. Ney, "Extensions of the sign language recognition and translation corpus rwth-phoenix-weather." in *LREC*, 2014, pp. 1911–1916.

[47] J. Huang, W. Zhou, Q. Zhang, H. Li, and W. Li, "Video-based sign language recognition without temporal segmentation," in *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[48] C. H. Lampert, H. Nickisch, and S. Harmeling, "Learning to detect unseen object classes by between-class attribute transfer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 951–958.

[49] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth, "Describing objects by their attributes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2009, pp. 1778–1785.

[50] V. Ferrari and A. Zisserman, "Learning visual attributes," in *Proc. Adv. Neural Inf. Process. Syst.*, 2008, pp. 433–440.

[51] D. Jayaraman and K. Grauman, "Zero-shot recognition with unreliable attributes," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3464–3472.

[52] D. Parikh and K. Grauman, "Relative attributes," in *Proc. IEEE Int. Conf. on Computer Vision*, 2011, pp. 503–510.

[53] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid, "Label-embedding for attribute-based classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2013, pp. 819–826.

[54] C. H. Lampert, H. Nickisch, and S. Harmeling, "Attribute-based classification for zero-shot visual object categorization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 3, pp. 453–465, 2014.

[55] Z. Li, E. Gavves, T. Mensink, and C. G. Snoek, "Attributes make sense on segmented objects," in *European Conference on Computer Vision*, 2014, pp. 350–365.

[56] G. Kulkarni, V. Premraj, V. Ordonez, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg, "Babytalk: Understanding and generating simple image descriptions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 12, pp. 2891–2903, 2013.

[57] V. Ordonez, G. Kulkarni, and T. L. Berg, "Im2text: Describing images using 1 million captioned photographs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2011, pp. 1143–1151.

[58] V. Escorcia, J. Carlos Niebles, and B. Ghanem, "On the relationship between visual attributes and convolutional networks," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1256–1264.

[59] P. Peng, Y. Tian, T. Xiang, Y. Wang, and T. Huang, "Joint learning of semantic and latent attributes," in *European Conference on Computer Vision*, 2016, pp. 336–353.

[60] M. Rohrbach, M. Stark, and B. Schiele, "Evaluating knowledge transfer and zero-shot learning in a large-scale setting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2011, pp. 1641–1648.

[61] M. Elhoseiny, B. Saleh, and A. Elgammal, "Write a classifier: Zero-shot learning using purely textual descriptions," in *Proc. IEEE Int. Conf. on Computer Vision*, 2013, pp. 2584–2591.

[62] T. Mensink, E. Gavves, and C. G. Snoek, "Costa: Co-occurrence statistics for zero-shot classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 2441–2448.

[63] J. Lei Ba, K. Swersky, S. Fidler *et al.*, "Predicting deep zero-shot convolutional neural networks using textual descriptions," in *Proc. IEEE Int. Conf. on Computer Vision*, 2015, pp. 4247–4255.

[64] Y. Fu, T. M. Hospedales, T. Xiang, Z. Fu, and S. Gong, "Transductive multi-view embedding for zero-shot recognition and annotation," in *Proc. European Conf. on Computer Vision*, 2014, pp. 584–599.

[65] B. Romera-Paredes and P. Torr, "An embarrassingly simple approach to zero-shot learning," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2152–2161.

[66] J. Qin, L. Liu, L. Shao, F. Shen, B. Ni, J. Chen, and Y. Wang, "Zero-shot action recognition with error-correcting output codes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 2833–2842.

[67] R. Qiao, L. Liu, C. Shen, and A. Van Den Hengel, "Less is more: zero-shot learning from online textual documents with noise suppression," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 2249–2257.

[68] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 3111–3119.

[69] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proc. of conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.

[70] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019, pp. 4171–4186.

[71] J. Liu, B. Kuipers, and S. Savarese, "Recognizing human actions by attributes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2011, pp. 3337–3344.

[72] M. Jain, J. C. van Gemert, T. Mensink, and C. G. Snoek, "Objects2action: Classifying and localizing actions without any video example," in *Proc. IEEE Int. Conf. on Computer Vision*, 2015, pp. 4588–4596.

[73] X. Xu, T. M. Hospedales, and S. Gong, "Semantic embedding space for zero-shot action recognition," *2015 IEEE International Conference on Image Processing (ICIP)*, pp. 63–67, 2015.

[74] X. Xu, T. Hospedales, and S. Gong, "Transductive zero-shot action recognition by word-vector embedding," *International Journal of Computer Vision*, vol. 123, no. 3, pp. 309–333, 2017.

[75] Q. Wang and K. Chen, "Alternative semantic representations for zero-shot human action recognition," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2017, pp. 87–102.

[76] A. Habibian, T. Mensink, and C. G. Snoek, "Video2vec embeddings recognize events when examples are scarce," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 10, pp. 2089–2103, 2017.

[77] M. Hahn, A. Silva, and J. M. Rehg, "Action2vec: A crossmodal embedding approach to action learning," in *The British Machine Vision Conference (BMVC)*, September 2018.

[78] W. Thomason and R. A. Knepper, "Recognizing unfamiliar gestures for human-robot interaction through zero-shot learning," in *International Symposium on Experimental Robotics*, 2016, pp. 841–852.

[79] N. Madapana and J. P. Wachs, "Hard zero shot learning for gesture recognition," in *IAPR International Conference on Pattern Recognition*, 2018, pp. 3574–3579.

[80] M. T. Ribeiro, S. Singh, and C. Guestrin, ""Why Should I Trust You?": Explaining the Predictions of Any Classifier," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Aug. 2016.

[81] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," in *NIPS*, 2017.

[82] H. Chen, S. Lundberg, and S.-I. Lee, "Explaining Models by Propagating Shapley Values of Local Components," in *AAAI*, 2019.

[83] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, "Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization," in *Proc. IEEE Int. Conf. on Computer Vision*, 2017.

[84] R. R. Selvaraju, P. Chattopadhyay, M. Elhoseiny, T. Sharma, D. Batra, D. Parikh, and S. Lee, "Choose Your Neuron: Incorporating Domain Knowledge Through Neuron-Importance," in *Proc. European Conf. on Computer Vision*, Cham, 2018, vol. 11217, pp. 540–556.

[85] Y. Liu and T. Tuytelaars, "A Deep Multi-Modal Explanation Model for Zero-Shot Learning," *IEEE Trans. on Image Processing*, vol. 29, pp. 4788–4803, 2020.

[86] E. Costello, *Random House Webster's Concise American Sign Language Dictionary*. Random House, 1999.

[87] R. A. Tennant, M. Gluszak, and M. G. Brown, *The American sign language handshape dictionary*. Gallaudet University Press, 1998.

[88] Y. Xian, B. Schiele, and Z. Akata, "Zero-shot learning-the good, the bad and the ugly," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 4582–4591.

[89] C. Lucas, R. Bayley, and C. Valli, *What's your sign for pizza?: An introduction to variation in American Sign Language*. Gallaudet University Press, 2003.

[90] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 6299–6308.

[91] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 1–9.

[92] A. Graves and J. Schmidhuber, "Framewise phoneme classification with bidirectional lstm and other neural network architectures," *Neural Networks*, vol. 18, no. 5-6, pp. 602–610, 2005.

[93] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[94] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *EMNLP*, 2014.

[95] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh, "Openpose: Realtime multi-person 2d pose estimation using part affinity fields," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

[96] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.

[97] J. Weston, S. Bengio, and N. Usunier, "Large scale image annotation: learning to rank with joint word-image embeddings," *Machine learning*, vol. 81, no. 1, pp. 21–35, 2010.

[98] G. Sumbul, R. G. Cinbis, and S. Aksoy, "Fine-grained object recognition and zero-shot learning in remote sensing imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 2, pp. 770–779, 2018.

[99] E. Kodirov, T. Xiang, and S. Gong, "Semantic autoencoder for zero-shot learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2017, pp. 3174–3183.

[100] V. Ferrari and A. Zisserman, "Learning visual attributes," *Proc. Adv. Neural Inf. Process. Syst.*, vol. 20, pp. 433–440, 2007.

[101] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[102] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata, "Feature generating networks for zero-shot learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 5542–5551.

[103] S. Narayan, A. Gupta, F. S. Khan, C. G. Snoek, and L. Shao, "Latent embedding feedback and discriminative features for zero-shot classification," in *ECCV*, 2020, pp. 479–495.

[104] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*, 2017, pp. 214–223.

[105] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.

Yunus Can Bilge received his BSc degrees from Izmir University of Economics, Turkey and received his MEngSc degree from University of New South Wales, Australia. He is a PhD candidate at Hacettepe University, Turkey. His research interests include computer vision and machine learning with minimal supervision.

Ramazan Gokberk Cinbis received his Bsc degree from Bilkent University, Turkey, and received his M.A. degree from Boston University, USA. He was a doctoral student at INRIA Grenoble between 2010-2014, and received a PhD degree from Université de Grenoble, France, in 2014. He is currently a faculty member at METU, Ankara, Turkey. His research interests include machine learning with limited and incomplete supervision.

Nazli Ikizler-Cinbis received her BSc and MSc degrees from Department of Computer Engineering at Bilkent University. During 2005-2006, she was a visiting scholar at University of Illinois at Urbana-Champaign (UIUC). After receiving her PhD degree from Bilkent University, she worked as a post-doctoral research associate at Boston University (USA). Since 2011, she is a faculty member at Hacettepe University. Her research areas are computer vision and machine learning, specifically focusing on video processing, human action and interaction recognition, and zero-shot learning.

APPENDIX

MS-ZSSLR-(W)ILD DATASET

MS-ASL [15] dataset contains significant variations due to various dialects in sign languages, and the sign language dictionaries do not comprehensively provide the dialect descriptions, to the best of our knowledge. Based on the MS-ASL [15] dataset, we have therefore created two separate benchmarks: MS-ZSSLR-W(ild) and MS-ZSSLR-C(lean). MS-ZSSLR-W contains the dialects, and MS-ZSSLR-C contains only the samples where the textual descriptions are inline with the visual descriptions of signs, as also explained in the main paper.

For our MS-ZSSLR-C/W benchmarks, we collect text and attribute based class definitions from Webster American Sign Language Dictionary [86] and American Sign Language Handshape Dictionary [87]. Figure 9 shows examples for demonstrating the aforementioned variances of dialects, and the problem of visual sign and definition mismatches. The first column shows an example video that is inline with the dictionary description. The remaining columns show different visual variations. We observe that the differences are mostly due to the variations in handshape, hand location with respect to the body and hand movements in sign making. Some signers execute the sign with two hands while the sign is supposed to be executed with one hand according to the textual definition or vice-versa. We also observe variations due to left-handed signers in the dataset.

ATTRIBUTES

Attribute-based representations are defined with respect to four main parameters: *handshape*, *palm orientation*, *hand location*, *hand movement*. There are two other attributes describing whether the sign includes repeated movements, and whether the sign is executed with a single hand. Below is the list of attributes defined for ASL-Text dataset:

- **Handshape:** `A, Open A, B, Baby 0, Bent B, V, Bent V, Open B, C, D, E, F, G, H, I, K, L, M, Flattened O, S, X, V, Y, 1, 3, 4, 5, Bent 5, 8, Open 8`.
- **Palm Orientation:** `in, out, up, down, left, right`.
- **Hand Location:** `neutral, chest, ear/temple, cheek/chin/mouth/nose, forehead/eyes, on shoulder`.
- **Hand Movement:** `move towards to the body, move away from the body, up, down, left, right, circular, internal movement at the wrist, internal movement at the fingers`.
- **Others:** `movement repeats?, one-hand sign?`.

As MS-ZSSLR-C/W datasets include different sign classes, handshape attributes are slightly altered to define new sign classes. The corresponding handshape attributes that are used in MS-ZSSLR-C/W datasets are as follows:

- **Handshape:** `A, Open A, B, Baby 0, Bent B, Open B, C, D, F, G, H, K, L, Bent L, Open N, N, Flattened O, Baby O, R, S, T, W, X, V , Y, 1, 5, Bent 5, 8, Open 8`.

DATASET SPLITS

Figure 10 shows average sample count per train, validation and test sets of ASL-Text, MS-ZSSLR-W and MS-ZSSLR-C datasets. We provide per-dataset split details in the following subsections.

ASL-Text Dataset

ASL-Text dataset contains 170, 30, and 50 mutually exclusive classes in train, validation and test sets. The class names are listed below.

Training classes: Excuse, Include/Involve, Shelf/Floor, Answer, Boss, Dress/Clothes, Marry, Stand-Up, Disappoint, Expert, Cancel/Criticize, Fed-Up/Full, Guitar, Emphasize, Government, Look, Afraid, Court, Medicine, Hello, Conflict/Intersection, Less, Of-Course, Dismiss, Dark, Silly, Home, Blue, Appointment, Disconnect, A-Lot, Enter, Mad, Cold, Decide, Arrive, Inform, Proceed, Miss/Assume, Letter/Mail, Keep, Fly-By-Plane, Deaf, High, Beautiful, Again, Happen, Depress, Embarrass, Deposit, Drunk, Develop, Over/After, Brave/Recover, Avoid/Fall-Behind, Full, Blame, Goal, Art/Design, Allow, Live, Future, Boy, Nice/Clean, Dry, Have, Take-Up, Heavy, Grow, Earth, Friday, Down, Bore, Center, Cheap, Everyday, Divorce, Forget, Awkward, Grandfather, Cruel, Graduate, Beer, Leave-There, Any, Chemistry, Brown, Friend, Left, Free, Freeze, Cannot, All, East, Give-Up, Family, Bad, Green, Can, Learn, Coat, Drink, Head, Football, Lousy, Buy, Excited, Price, Enough, Grandmother, Lie, Sausage/Hot-Dog, Thrill/Whats-Up, Shame, Same-Old, Messed-Up, Match, Bar, Car, Helmet, Illegal, Merge/Mainstream, Chase, Work-Out, Weekend, Dirty, How-Many/Many, Gone, Far, Head-Cold, Chain/Olympics, Line, Go-Away, Collect, Set-Up, Country, Really, Protest, Flat-Tire, Lungs, Paint, Inject, Easy, Lip/Mouth, Nab, Fail, Fence, To-Fool, Gamble, Banana, Introduce, Mosquito, Lend, Finally, Halloween, Exact, Hearing-Aid, Explain, Lecture, Bicycle, Magazine, Increase, Disappear, Make, Lose-Competition, Experience, Expensive, Girl, Accept, But.

Validation classes: Meet, Finish, Advise/Influence, Course, Destroy, Cough, Alone, Bridge, Call-By-Phone, Hard, Idea, Apple, Hospital, One-Month, Black, Grass, Borrow, Run-Out, Bread, Monday, Library, One, Metal, Morning, Hit, Most, Meat, Come-On, Not-Mind, Smooth.

Test classes: Generation, Half, Engagement, Break-Down, Apply, Date/Dessert, Shape/Statue, Pass, Insult, Hamburger, Obscure, Like, Crush, Less-Than, Bawl-Out, Blind, Paper-Check/Card, Get-Up, Place, Cooperate/Unite, Insurance/Infection, Follow, Meeting, General, Autumn, Comb, Experiment, Line-Up, Gas/Gas-Up, Grab-Chance, Permit, Eat, Tough, Trash/Bag, Speech/Oral, Cherish, Strange, Association, Pull, Member, Ghost, Machine, Average, Act, Ahead, Celebrate, Skin,Strong, Where, Concern.

MS-ZSSLR-C/W Dataset Zero-Shot Settings

MS-ZSSLR-C/W datasets contain 120, 30, and 50 mutually exclusive classes in train, validation and test sets, respectively. The corresponding class names are listed below.

Train classes: Again, Apple, Aunt, Bad, Bathroom, Beautiful, Bird, Black, Blue, Book, Boring, Boy, Bread, Brother,

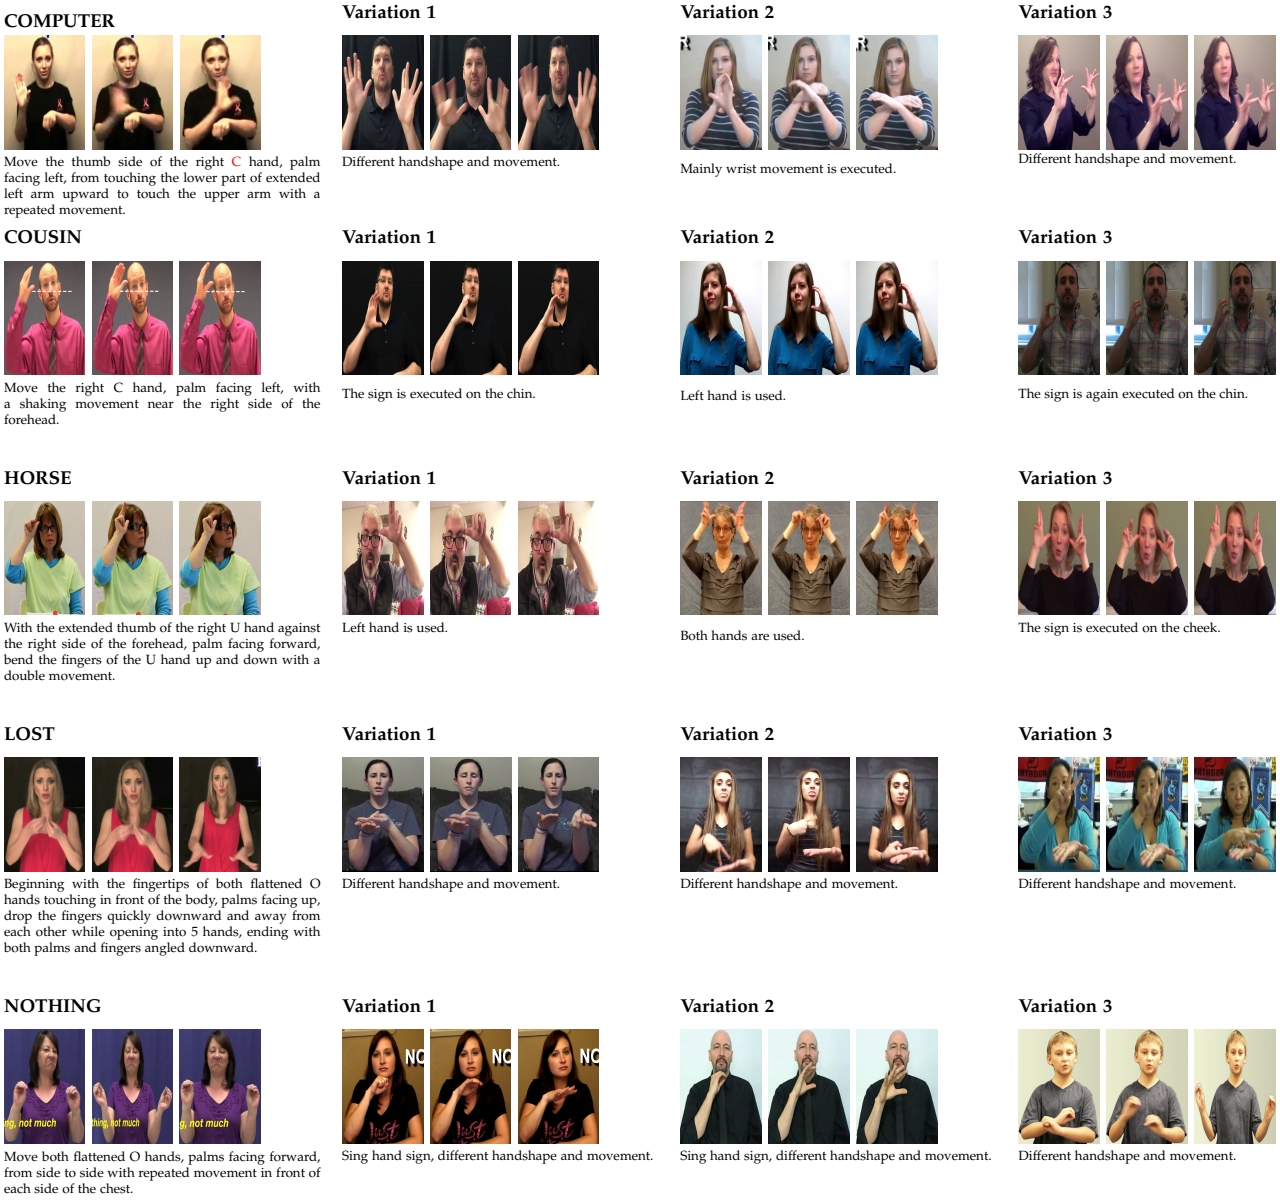

Fig. 9. Example in-class variations and corresponding textual descriptions from MS-ZSSLR-W(ild) dataset. The first column shows an example video that is inline with the dictionary description. The remaining columns show different visual variations.

Brown, But, Call, Can, Cat, Cheese, Coffee, College, Color, Cook, Dance, Daughter, Day, Doctor, Draw, Drink, Eat, Family, Father, Fine, Finish, Fish, Forget, France, Friend, Future, Girl, Go, Good, Grandfather, Grandmother, Green, Have, Hearing, Hello, Help, Home, Horse, How-Many, Hungry, Hurt, Jacket, Know, Learn, Like, Man, Me, Milk, Mother, Must, Name, Nice, Night, No, Not, Nothing, Now, Nurse, Old, Orange, Paper, Pencil, Pink, Play, Please, Purple, Read, Red, Room, Sad, School, See, Sell, Shirt, Shoes, Sick, Sister, Sit, Slow, Sorry, Spring, Student, Sunday, Table, Teach, Teacher, Thank-You, Time, Tired, Tomorrow, Understand, Walk, Want, Water, What, When, Where, White, Who, Woman, Work, Write, Wrong, Yellow, Yes, You.

Validation classes: Afraid, Bicycle, Buy, Class, Cousin, Door, English, Hot, Late, Lost, Mad, Meet, Movie, Nephew, Not-Know, Not-Like, Remember, Restaurant, Run, Same, Sign, Son, Start, Tea, Week, Which, Why, Wife, Year, Your.

Test classes: All, Baby, Bed, Big, Boyfriend, Candy, Car, Cold, Computer, Cookie, Deaf, Dentist, Divorce, Dog, Early, Enjoy, Excited, Favorite, Germany, Happy, Here, Hotdog, Homework, Hour, House, How, Kitchen, Library, Live, Love, Mean, Money, More, My, Party, Practice, Ready, Right, Sandwich, Saturday, Soccer, Soda, Study, Today, Turkey, Ugly, Uncle, Watch, Wednesday, Yesterday.

Additional Qualitative Results

Figure 11 and 12 present examples for correct and incorrect classifications. We observe that most confusions occur across classes with similar hand movements and/or locations, highlighting the importance of handshape pattern modeling. We also observe that class definitions also tend

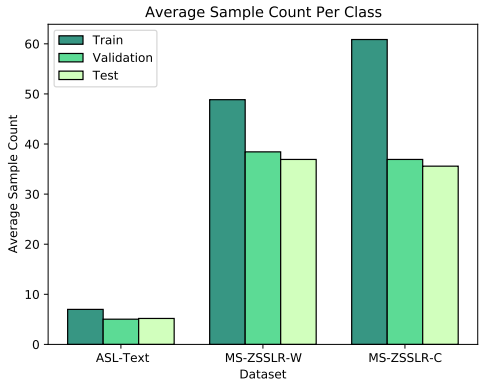

Fig. 10. Average sample count per train, validation and test sets for ASL-Text, MS-ZSSLR-C, and MS-ZSSLR-W datasets.

MS-ZSSLR-C

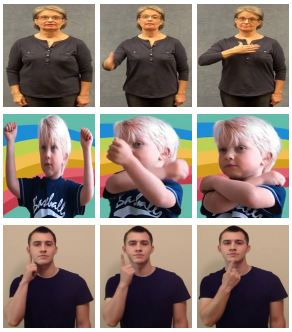

MY
Place the palm of the right open hand on the chest, fingers pointing left.

LOVE
With the wrists of both S hands crossed in front of the chest, palms facing in, bring the arms back against the chest.

DEAF
Touch the extended right index finger first to near the right ear and then to near the right side of the mouth.

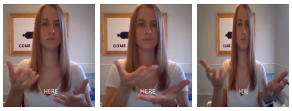

HERE
Beginning with both curved hands in front of each side of the body, palms facing up, move the hands toward each other in repeated flat circles.

HAPPY
Brush the fingers of the right open hand, palm facing in and fingers pointing left, upward in a repeated circular movement on the chest.

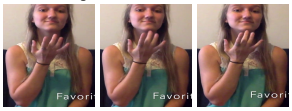

FAVORITE
Touch the bent middle finger of the right 5 hand, palm facing in, to the chin with a double movement.

SANDWICH
With the palms of both open hands together, right hand above left, bring the fingers back toward the mouth with a short double movement.

BIG
Move both L hands from in front of each side of the chest, palms facing each other, in large arcs beyond each side of the body.

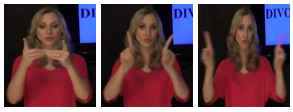

DIVORCE
Beginning with the fingertips of both the D hands touching in front of chest, palms facing each other and index fingers pointing up, swing the hands away from each other by twisting the wrists, ending with the hands in front of each side of the body, palms facing forward.

Fig. 11. Example predictions of the proposed model on the MS-ZSSLR-C dataset. The first three rows show examples that are correctly predicted and the last three rows show incorrect predictions. For the correct predictions, second column includes corresponding textual description. For the incorrect predictions, the second column includes ground truth textual descriptions and frames from an example video of the ground truth class.

ASL-Text

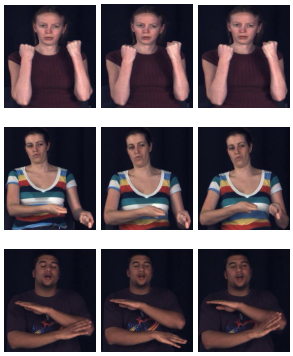

STRONG
Move both S hands, palms facing in, forward with a short deliberate movement from in front of each shoulder.

PULL
Beginning with the right curved hand in front of the body and the left curved hand somewhat forward, both palms facing up, bring the hands back toward the right side of the body while closing them into A hands.

AUTUMN
Brush the index-finger side of the right B hand, palm facing down, downward toward the elbow of the left forearm, held bent across the chest.

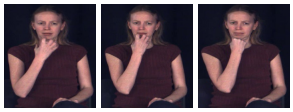

CHERISH
Beginning with the right curved 5 hand in front of the mouth, palm facing back, slowly close the fingers into an S hand.

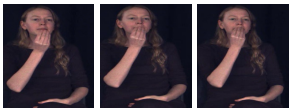

EAT
Bring the fingertips of the right flattened O hand, palm facing in, to the lips with a repeated movement.

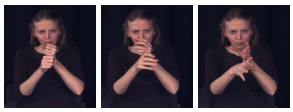

BAWL-OUT
Beginning with the little finger of the right S hand on the top of the index finger side of the left S hand, flick the hands forward with a deliberate double movement while opening the fingers into 5 hands each time.

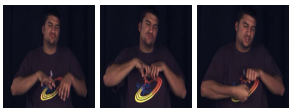

CRUSH
Beginning with the palms of both A hands together in front of the chest, twist the hands in opposite directions.

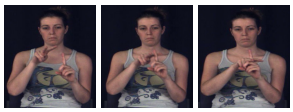

APPLY
Move the fingers of the right V hand, palm facing forward, downward on each side of the extended left index finger, pointing up in front of the chest.

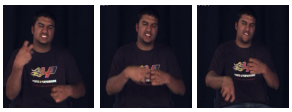

TOUGH
Beginning with both bent V hands in front of the chest, right hand higher than the left hand, palms facing in, move the right hand down and the left hand upward with an alternating movement, brushing the knuckles of each hand as the hands move in the opposite direction.

Fig. 12. Example predictions of the proposed model on the ASL-Text dataset. The first three rows show examples that are correctly predicted and the last three rows show incorrect predictions. For the correct predictions, second column includes corresponding textual description. For the incorrect predictions, the second column includes ground truth textual descriptions and frames from an example video of the ground truth class.

to be similar across the confused classes. These qualitative examples further demonstrate that the problem domain can benefit from more detailed visual and auxiliary data representations.