

Multi-Label Zero-Shot Human Action Recognition via Joint Latent Ranking Embedding

Qian Wang, Ke Chen

The University of Manchester, UK.
qian.wang173@hotmail.com, ke.chen@manchester.ac.uk

Abstract

Human action recognition refers to automatic recognizing human actions from a video clip, which is one of the most challenging tasks in computer vision. Due to the fact that annotating video data is laborious and time-consuming, most of the existing works in human action recognition are limited to a number of small scale benchmark datasets where there are a small number of video clips associated with only a few human actions and a video clip often contains only a single action. In reality, however, there often exist multiple human actions in a video stream. Such a video stream is often weakly-annotated with a set of relevant human action labels at a global level rather than assigning each label to a specific video episode corresponding to a single action, which leads to a multi-label learning problem. Furthermore, there are a great number of meaningful human actions in reality but it would be extremely difficult, if not impossible, to collect/annotate video clips regarding all of various human actions, which leads to a zero-shot learning scenario. To the best of our knowledge, there is no work that has addressed all the above issues together in human action recognition. In this paper, we formulate a real-world human action recognition task as a multi-label zero-shot learning problem and propose a framework to tackle this problem in a holistic way. Our framework holistically tackles the issue of unknown temporal boundaries between different actions for multi-label learning and exploits the side information regarding the semantic relationship between different human actions for knowledge transfer. As a result, our framework leads to a joint latent ranking embedding for multi-label zero-shot human action recognition. A novel neural architecture of two component models and an alternate learning algorithm are proposed to carry out the joint latent ranking embedding learning. Thus, multi-label zero-shot recognition is done by measuring relatedness scores of action labels to a test video clip in the joint latent visual and semantic embedding spaces. We evaluate our framework with different settings, including a novel data split scheme designed especially for evaluating multi-label zero-shot learning, on two weakly annotated multi-label human action datasets: *Breakfast* and *Charades*. The experimental results demonstrate the effectiveness of our framework in multi-label zero-shot human action recognition.

Keywords: Human action recognition, Multi-label learning, Zero-shot learning, Joint latent ranking embedding, Weakly supervised learning

1. Introduction

As one of the most challenging tasks in computer vision, human action recognition refers to automatic recognizing human actions conveyed in a video clip. In last two decades, human action recognition has been extensively studied. As there are many different human actions in reality, this task is generally formulated as a multi-class classification problem. To train a multi-class classifier for human action recognition, a great number of examples for each single action are required in the current setting. To collect such training examples, one needs to manually trim a video stream to ensure that there is only one human action appearing in a trimmed video episode. This annotation process is laborious and time-consuming and there is hence no large-scale dataset with “fine-grained” annotation for human action recognition. In contrast to ImageNet [1] for object recognition, where it consists of a total of 3.2 million cleanly labelled images spreading over 5,247 categories, there are much fewer annotated video clips involving only a small number of human actions. For instance, HMDB51 and UCF101 are among the most commonly used benchmark datasets in human action

recognition, where there are 6,676 and 13,320 instances of only 51 and 101 different human actions, respectively. The limitation of human action datasets in such a scale has become an obstacle in developing a large-scale human action recognition system.

In a real scenario, a video clip often conveys multiple human actions corresponding to different concepts. Hence, a set of multiple action labels have to be used to characterize its complete semantics underlying human actions conveyed in this video clip. For example, video clips on YouTube are usually uploaded by users along with some descriptive terms that can be used to infer the human actions conveyed in those video clips. In this circumstance, descriptive terms may be viewed as a set of coherent labels that collectively characterize the semantics at a global level. Recently, a very large multi-label video dataset YouTube-8M [2] has been collected by Google Research. Although the dataset is not restricted to human action video clips, it paves a new way for various video analyses including human action recognition. One of essential video analysis problems on such a data set may be formulated as multi-label learning that predicts a set of labels associated to a given instance or a

set of relatedness scores corresponding to the candidate labels that could characterize this instance. In multi-label learning, a training example often consists of an input instance and a set of labels associated with this instance at a global level (no need of explicitly associating each of those labels to a relevant object within this instance). While multi-label data are common in many domains and multi-label classification has been studied under different applications [3], e.g., semantic image tagging, text categorization and gene functionality prediction, only few works are pertinent to multi-label human action recognition in literature due to a lack of human action datasets annotated with multiple class labels. To fill in this gap, a dataset dubbed Charades [4] was collected especially for multi-label human action recognition and made publicly available very recently. In addition, other datasets collected for different tasks were also considered to be used in multi-label human action recognition. Thus, such data sets provide a proper test bed for multi-label human action recognition studies.

Multi-label human action recognition often has to work on weakly labelled video data, i.e., the training data are annotated at the video level without exhaustively trimming and annotating multiple action episodes. While it is easier to collect such video clips associated with a set of labels at a global level than those with “fine-grained” annotation, it would be still very challenging to collect all the training examples due to the existence of many different human actions. *Zero-shot learning* (ZSL) provides an alternative solution to alleviate this problem. ZSL aims to recognize the instances belonging to novel classes which are not seen during training. It has been formally shown that under certain conditions, a ZSL system trained on a dataset of finite classes could be used to predict infinite number of classes unseen during the learning [5]. Under the ZSL framework, we merely need to collect and annotate training examples for a moderate set of training classes and expect that a large number of novel classes can be recognized via exploiting the semantic relationship between different human actions. To this end, a ZSL algorithm needs to transfer the knowledge regarding the relations between visual features and class label semantics learned from known or training classes to unseen or test classes. The knowledge transfer is enabled by modelling the *semantic representations* of different classes, which can be easily obtained from side information, e.g., descriptive texts, with a much less effort than collecting and annotating visual data. Nevertheless, most of the existing ZSL methods were proposed to tackle single-label ZSL problems but multi-label ZSL problems are much more complicated, leading to additional challenge that do not exist in single-label ZSL. Although some of single-label ZSL methods might be extensible to multi-label scenarios, their effectiveness of different ZSL algorithms have not been extensively investigated in the multi-label learning scenarios. To the best of our knowledge, there exists no work in multi-label zero-shot human action recognition.

In this paper, we address the multi-label ZSL issues in the context of human action recognition. In our problem, the training video data are weakly labelled so that the exact temporal-spatial locations of multiple human actions in a video clip remain unknown. In addition, training examples consisting of vi-

sual instances and their corresponding label sets of multiple labels are only available for those associated with training/known labels, a subset of the action label collection considered in the recognition stage. Thus, the nature of multi-label zero-shot human action recognition poses several challenges that do not exist in static data and single-label ZSL. To tackle all the challenges in a holistic way, we propose a novel *joint latent ranking embedding* framework. The framework aims to learn a joint latent ranking embedding from visual and semantic domains. By using the learned joint latent ranking embedding, any visual instances and any action labels can be mapped into the joint latent visual and semantic embedding spaces where positive connections between visual instances and action labels rank ahead of negative ones. Thus, any human actions can be recognized regardless of known or unseen actions during learning. Our framework consists of two component models: *visual* and *semantic* models. The visual model learns mapping a visual instance into the latent *visual embedding* space, while the semantic model learns mapping action labels into the latent *semantic embedding* space. The visual and the semantic models are tightly coupled to learn a proper ranking that works in the joint latent visual and semantic embedding spaces with an alternate learning algorithm on training examples annotated with only known action labels. In the test, multi-label zero-shot recognition is done by measuring relatedness scores of action labels to a test visual clip in the joint latent visual and semantic embedding spaces.

Our main contributions in this paper are summarized as follows:

- By considering real scenarios, we formulate general human action recognition as a *multi-label zero-shot learning* problem. To the best of our knowledge, our work presented in this paper is the first attempt in studying human action recognition from a multi-label zero-shot learning perspective, which tackles several technical challenges pertaining to this problem in a holistic way.
- To address the multi-label zero-shot issues arising from weakly annotated data for human action recognition, we propose a novel joint latent ranking embedding framework consisting of visual and semantic embedding models. To train two embedding models effectively, we come up with a learning algorithm that alternately optimizes the parameters in two embedding models via minimizing the proper rank loss functions.
- To test the performance of our proposed framework, we conduct a thorough evaluation via a comparative study on two benchmark multi-label human action datasets, *Breakfast* and *Charades*, with various evaluation metrics and different settings including a novel data split protocol simulating a real scenario of multi-label zero-shot human action recognition.

The rest of this paper is organized as follows. Section 2 reviews related works. Section 3 presents our framework for multi-label zero-shot human action recognition. Section 4 de-

scribes our experimental settings, and Section 5 reports the experimental results. The last section draws conclusions.

2. Related Work

In this section, we review the existing works relating to multi-label human action recognition, especially for those applicable to multi-label zero-shot learning scenarios, and point out the limitations of the existing works. We first overview existing multi-label classification methods and then focus on the existing works in multi-label ZSL learning despite the fact that none of such multi-label ZSL methods has been applied to human action recognition. Finally, semantic representations required by any ZSL methods are briefly reviewed.

2.1. Multi-label Learning

In a real scenario, semantics underlying real-world data is often complex and has to be characterized with multiple labels, e.g., web videos. In a video clip pertaining to human actions, multiple actions could happen simultaneously, e.g., sitting, eating and listening. In this scenario, no episode in such a video clip can be characterized by a single action label and a set of labels hence have to be used collectively to describe this video clip. Even though a video clip can be divided into several episodes corresponding to different human actions, the segmentation and annotation process could be difficult, tedious, laborious and time-consuming. In particular, semantic image segmentation and human action detection in video streams remain unsolved in general. As a result, multi-label learning is often formulated as a weakly supervised learning task that predicts a set of labels associated with an instance but does not address the issue in assigning each label in the set to a specific object within this instance. To tackle a weakly supervised multi-label learning problem, two different representation methods are used to characterize input data: *instance-level* and *object-level* representations. An instance-level representation is a global representation of an instance, e.g., a video clip or an image, without considering objects appearing in this instance, while an object-level representation is a local representation that describes individual objects extracted from an instance, e.g., semantically meaningful episodes/patches in the video/image. Depending on the representation of input data, multi-label learning methods can be divided into two categories.

In multi-label learning, most of the existing methods [6, 7, 8, 9] work on an instance-based representation, a single feature vector of an instance. Recently, Fast0Tag [6] was proposed for multi-label image tagging by learning a mapping from visual to label space. An image containing multiple objects is represented by one aggregated visual representation. Alternatively, TagProp [7] uses an adapted nearest neighbour model for multi-label learning in visual space where each image of multiple objects is also represented by one feature vector at the instance level. [8] use a convolutional neural network (CNN) directly working on raw images of multiple objects to learn image-level deep visual representations for multi-label classification. [9] use a deep neural network with a rank loss in learning for large-scale multi-label text classification where an input document is

represented with a single feature vector. Although representing one instance at the global level is straightforward and convenient, it might neglect the intrinsic relationship between multiple objects within an instance. Thus, an instance-level representation might result in a catastrophic information loss, especially for long-term dependent and complex video data.

To overcome the weakness in neglecting the information regarding the intrinsic relationship between objects within an instance, efforts have been made to exploit such information in previous works. Despite being difficult, the segmentation of multiple objects within an instance turns out to be beneficial to multi-label learning. One framework named *multi-instance multi-label learning (MIML)* [10] demonstrates that multi-label learning can be fulfilled effectively if multiple objects within an instance have been explicitly separated or segmented even if no label is explicitly assigned to each of multiple objects within an instance during learning. In real applications, however, automatic semantic segmentation of objects in an instance is also challenging, and a manual segmentation process is laborious and time-consuming. Moreover, some recent works tend to explore object-based representations without using any explicit semantic object segmentation techniques, which seeks a synergy between the MIML and object-level representations. [11] address this weakly supervised issue in multi-label human action detection with a two-stage solution. First, a set of potential objects or spatial-temporal volumes are generated and selected from a video instance with a set of handcrafted rules. Then the problem is transformed into a MIML problem which can be solved by those traditional multi-label learning algorithms under the MIML framework. Similar ideas were also explored by [12] and [13] for multi-label image classification. However, the extraction of true positive objects from the original visual instance is a very challenging yet non-trivial task, which critically determines the multi-label learning performance. To extract all the meaningful objects within an instance, a lot of candidate proposals have to be considered so that it might suffer from a high computational burden. Instead of using the MIML, [14] attempt to explore the information regarding multiple objects in instances via a matrix completion method. Their method works on the assumption that an instance representation may be expressed by a linear combination of hidden representations of objects appearing in this instance. Experimental results reported by [14] demonstrate the effectiveness of this method via an instance-level bag-of-words image representation. However, this idea does not seem applicable to other kinds of visual representations, such as those popular yet powerful deep representations.

2.2. Multi-label Zero-shot Learning

Zero-shot learning (ZSL) has attracted much attention in recent years and provides a promising technique for recognizing a large number of classes without the need of the training data concerning all the classes. Very recently, [5] have formally shown that it is feasible to predict a collection of infinite unseen labels with a classifier learned on training data concerning only a number of labels in this collection or a subset of this

collection, where multi-label ZSL is a special case in this so-called “infinite-label learning” paradigm. According to a ZSL taxonomy [15], existing ZSL approaches are divided into three categories, namely, *direct mapping* [16, 17, 18, 19, 20], *model parameter transfer* [21, 22] and *joint latent space learning* [23, 24, 25, 26, 27, 28, 15, 29]. Although most existing works focus on single-label ZSL, efforts have been made to address more complex multi-label ZSL issues [22, 6, 30, 31, 32, 33].

For *direct mapping*, it needs to learn a mapping directly from visual to semantic space for zero-shot recognition on the semantic space, which poses a challenge to multi-label ZSL. In single-label ZSL, a training example provides a visual-semantic representation pair used to learn a one-to-one direct mapping. In multi-label ZSL, however, one instance has to be associated with a set of multiple labels and the number of labels associated with different instances are various. As a label is represented with a semantic feature vector, e.g., a vector of attributes or a word vector, in a semantic space, it is no longer straightforward to learn a direct mapping from visual to semantic space in the context of multi-label ZSL. How to model complex semantics underlying a set of labels associated with an instance becomes a central issue in multi-label ZSL. To tackle this issue, most of existing works [33, 31] make use of the composition properties of semantic representations such as word vectors by using the average of semantic representations of multiple labels to a collective semantic representation for a set of labels associated with the instance. Thus, a training example is formed with a pair of an instance-level visual representation and its corresponding collective semantic representation, which enables one to learning a direct mapping for multi-label ZSL. Apparently, such a collective representation cannot avoid information loss even though a contextualized semantic representation [31] was used. In particular, the multi-label ZSL method proposed by [33] is subject to a fundamental limitation; their method has to take into account all the possible combinations of different unseen labels in a pre-fixed unseen label collection. Thus, the computational complexity of their algorithm grows exponentially with respect to the number of unseen labels and hence can cope with only a very small number of pre-fixed unseen labels (e.g., up to eight in their experiments). To alleviate the information loss problem in generating a collective semantic representation, Fast0Tag [6] introduces an alternative solution to collective semantic representations. In their method, each visual instance is mapped into a “principal direction” in the semantic space based on an assumption that there is always such a direction for any multi-labelled instances in a semantic space, e.g., word vector space, and all the labels associated with this instance always rank ahead of irrelevant labels. In other words, a hyperplane perpendicular to this direction can always be found to separate the relevant labels from the irrelevant ones for any multi-labelled instance. While this assumption holds for those datasets used in their zero-shot image tagging experiments [6], it remains unclear for other image datasets and different domains, e.g., human action recognition. From an alternative perspective, [30] suggest using an object-level visual presentation under the direct mapping framework for multi-label zero-shot object recognition. Before multi-label learning takes place, an

image thus has to be semantically segmented into meaningful subregions and each subregion can be characterized by one label. As a result, their solution is actually a special case of the MIML [10] but heavily relies on sophisticated semantic segmentation techniques that remain unavailable up to date. Furthermore, this method is not extensible to sequential data such as video clips.

Like the works in extending *direct mapping* to multi-label ZSL, the *model parameter transfer* idea is also adapted for multi-label ZSL, leading to COSTA [22]. COSTA aims to establish a model for each unseen label via a linear weighted combination of known-label models. The known-label models are trained independently by means of a one-vs-rest binary classifier, e.g., support vector machines (SVMs). The combination coefficients are determined by the co-occurrence of multiple labels derived from either annotation of datasets in hand or external web sources. In COSTA, however, the known-label models are trained independently without considering the relationship and coherence among those labels that together describe an instance. Then, COSTA only uses label co-occurrences to model the relatedness between a pair of labels but neglects the semantics of an individual label itself. So far, this idea has been tested only on static images in the context of multi-label zero-shot object recognition.

The *joint latent space learning* methodology was proposed for multimedia information retrieval and multi-label related learning and led to favorable results in real-world applications [34, 35, 36, 37]. The core idea underlying this methodology is learning a joint latent embedding from both visual and semantic domains to narrow the semantic gap so that a task can be done effectively in the latent embedding space(s). More recently, this general idea has also been explored in single-label ZSL [23, 24, 25, 26, 27, 28, 15, 29]. Empirical studies suggest that those *joint latent space learning* methods often outperform most of existing direct mapping and model parameter transfer methods on several benchmark datasets designed for single-label ZSL [23, 24, 25, 26, 27, 28, 15, 38]. To the best of our knowledge, however, there exists no joint latent space learning method to tackle multi-label ZSL problems probably due to a lack of techniques in modelling complex semantics underlying a set of multiple labels describing an instance, as elucidated above for direct mapping. In this paper, we propose a novel approach to multi-label zero-shot human action recognition by exploring the joint latent space learning idea to holistically tackle those challenges described in Section 1.

2.3. Semantic Representation

Regardless of different ZSL scenarios, modelling semantics underlying a collection of labels and their relatedness plays a critical role in knowledge transfer required by ZSL. Miscellaneous methods in semantics modelling and representations have been developed from different perspectives including *attributes of labels*, *label embedding*, *co-occurrence of labels* and *concept embedding*.

Attributes of labels are a generic semantic representation where a label is characterized by a list of attributes common to all the labels [16]. Label embedding refers to embedding

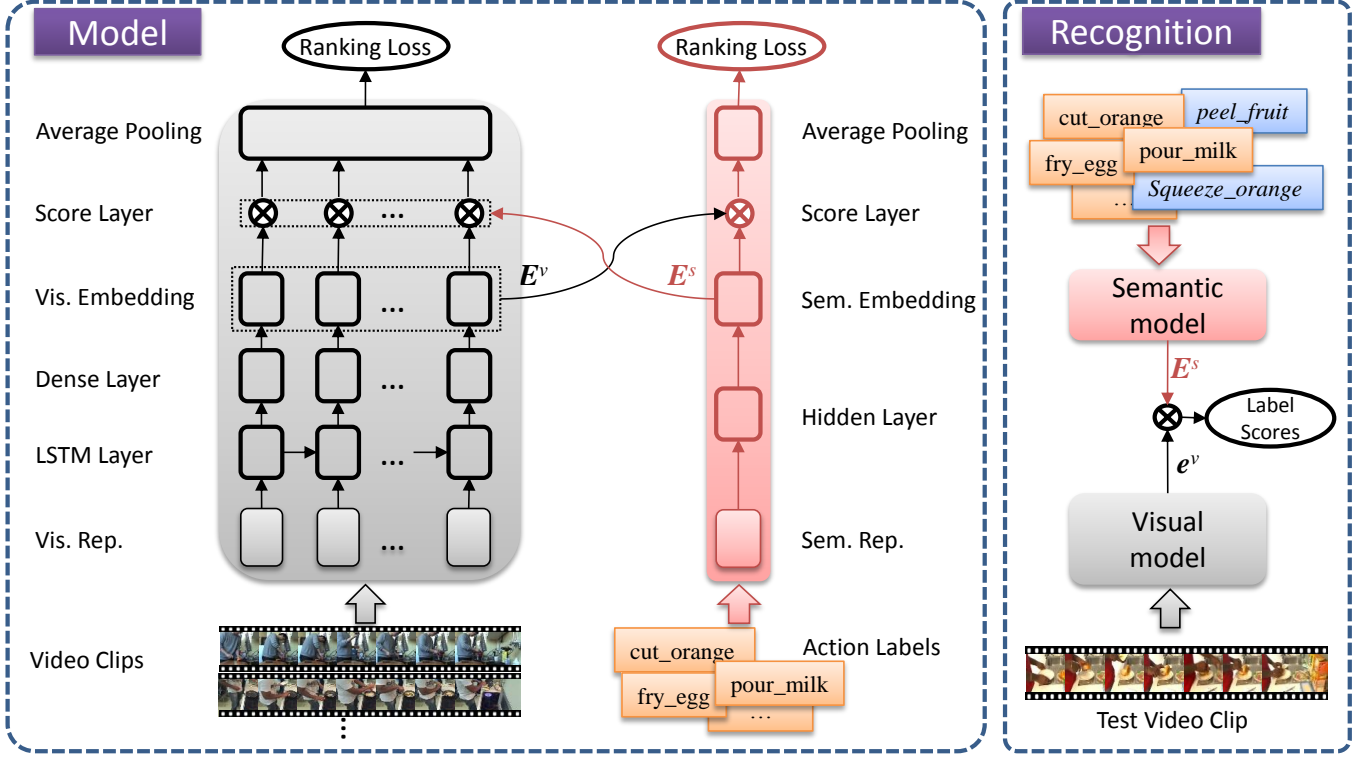


Figure 1: Our multi-label zero-shot human action recognition framework. This framework shown in the left box is composed of two component models: visual and semantic model highlighted with grey and red colors. After the joint latent ranking embedding learning, the trained visual and semantic models work together for multi-label zero-shot recognition, as shown in the right box. Action labels marked with brown color are training classes or known labels during learning, while action labels marked with blue colour are test classes or unseen labels during learning.

labels onto a semantic space where the semantic relatedness of labels are modeled [39]. Label embedding is often carried out via learning on external textual resources. For example, the famous *Word2Vec* semantic embedding is obtained by training a skip-gram neural network on the large-scale corpora, e.g., Google News dataset [39]. Such semantic representations are widely used in ZSL, e.g., [5, 6, 30, 33, 15]. Unlike label embedding obtained with external resources, co-occurrence of labels is yet another way to capture the relatedness between different labels, e.g., [32]. Alternatively, the co-occurrence information on different class labels can also be extracted from external resources for ZSL [22]. In particular, co-occurrence of labels allows for capturing the relatedness between labels jointly used to describe an instance. The label co-occurrence information may be incorporated into learning semantic embedding for a given dataset, e.g., [32]. In concept embedding, the semantic meaning of a label is assumed to be polysemous depending on different labels (together treated its context of this target label) jointly used to describe an instance. Hence, the semantic meaning of a label under a specific context frames a concept. As a result, concept embedding [40] can be viewed as contextualized label embedding where a label may have multiple semantic representations in different contexts. The concept embedding seems specific and is only applicable to direct mapping for multi-label ZSL [31]. Our proposed framework for multi-label zero-shot human action action is generic so that all the semantic represen-

tations apart from the concept embedding may be used directly in our framework.

3. Model Description

In this section, we present a novel framework for multi-label zero-shot human action recognition. First, we overview the proposed framework along with our motivation and justification. To make it self-contained, we then briefly review the LSTM unit, an important mechanism used in our framework for implicit saliency detection on video data. Next, we present the joint latent ranking embedding learning method including the rank loss functions and an alternate learning algorithm especially developed for our proposed architecture. Finally, we specify a procedure on how to apply a trained joint latent ranking embedding model to multi-label zero-shot human action recognition in test.

3.1. Overview

Our proposed framework aims at multi-label zero-shot human action recognition. We formulate this problem as learning a mapping $\phi: \mathbf{x} \rightarrow \mathbf{y}$, where \mathbf{x} is a visual input, e.g., a set of segment-level visual feature vectors extracted from a video clip, and $\mathbf{y} \in \mathbb{R}^{|C|}$ is a list of label-relatedness scores for \mathbf{x} with respect to a action label collection, $C = \{1, \dots, |C|\}$, where C is further divided into two mutually exclusive label subsets, C^{Tr}

and C^U , corresponding to known (training) and unseen actions; i.e., $C^{Tr} \cup C^U = C$ and $C^{Tr} \cap C^U = \emptyset$. During learning the mapping ϕ , only training examples of labels in C^{Tr} are available. However, the learned mapping ϕ is used to predict any actions appearing in a video clip no matter whether they are known actions in C^{Tr} or unseen actions in C^U .

To tackle the problem formulated above, we propose a joint latent ranking embedding framework. Motivated by the joint latent space learning idea used in ZSL [23, 24, 26, 27, 28, 15], we would tackle the knowledge transfer issue in the joint latent embedding spaces where the original visual and semantic representations are mapped into. By embedding visual and semantic representations into the joint latent embedding spaces, we expect that semantic gap can be narrowed considerably and the semantic relatedness of known and unseen labels may be effectively explored and exploited in zero-shot recognition. Thus, our framework consists of two component models: *visual* and *semantic* models used to learn latent visual and semantic embedding, respectively. Two component models are tightly coupled to learn a joint latent ranking embedding for knowledge transfer, as illustrated in the left box of Figure 1.

For visual embedding, we encounter two major technical issues due to the nature of weakly annotated data: a) for a visual input, it remains unknown where an episode conveying an action, and b) it remains unclear which of those action labels describing a video clip is associated with a specific video episode. Nevertheless, a video clip is an ordered sequence of frames and we could explore the temporal coherence underlying a video clip to tackle two aforementioned technical issues. Motivated by recent works in video classification and activity recognition [41, 42, 43], we employ a *long short-term memory* (LSTM) [44] recurrent neural network layer to capture temporal coherence underlying an action episode. Thus, the LSTM layer (c.f. Section 3.2) is first used to process a sequence of visual representations extracted from video segments. With the memorizing and forgetting mechanism of LSTM units, we expect that the LSTM layer explores the temporal structure of human actions conveyed in a video sequence; the LSTM units would memorize the input segments until parsing an episode regarding a human action is completed and then forget all the previous input segments when an episode conveying another action starts. Thus, an implicit saliency detection is carried out where no action episode boundaries are explicitly specified. For visual embedding, we further employ two fully-connected layers, dense layer of *rectified linear* (ReLU) units [45] and visual embedding layer of linear units, to capture salient features on the temporal coherence representations yielded by the LSTM layer. While this specific visual model shown in the left box of Figure 1 is used in our experiments, its capacity can be increased by adding more hidden units and/or layers if necessary. The score and average pooling layers above the visual embedding layer are used for joint latent ranking embedding learning as presented in Section 3.3. Thus, the visual model is carried out by a deep network of heterogeneous layers.

For semantic embedding, we employ a three-layer fully-connected neural network, the input layer, the hidden layer of ReLU units and the semantic embedding layer of linear units,

to carry out the semantic model, as shown in the left box of Figure 1. This learning model is capable of capturing the intricate semantic relatedness between different actions in a label collection of a moderate size, e.g. those datasets used in our experiments. If necessary, its capacity can be increased by adding more hidden units and/or layers. As a result, the neural network is fed with a specific semantic representation of action labels, e.g., word vectors and subsequently map them into the semantic embedding layer via a hidden layer. Likewise, the score and average pooling layers above the semantic embedding layer are used for joint latent ranking embedding learning. To explore the semantic relatedness between different labels in bridging the semantic gap between visual and semantic space, semantic embedding learning needs to automatically exploit the information carried in training data, e.g., frequency of label co-occurrence in a training dataset.

During the joint latent ranking embedding learning, the visual and semantic models are tightly coupled to learn a ranking criterion for the joint latent visual and semantic embedding spaces. This ranking criterion ensures that the relatedness scores of those labels associated with a visual instance are higher than those for other labels irrelevant to this instance, and the relatedness scores of those visual instances relevant to an action label are higher than those of other visual instances irrelevant to this label. For learning, we propose an algorithm working alternately on two models for parameter estimation by promoting the correct ranking based on training examples in known classes. During training, the visual model learns the visual embedding of a video instance such that those labels relevant to this instance rank ahead of other irrelevant ones in terms of the relatedness scores estimated on the semantic latent space, E^s . Reciprocally, the semantic model learns the semantic embedding of action labels such that the relevant visual instances rank higher than the irrelevant ones in terms of relatedness scores calculated in the visual latent space, E^v . Once the learning is completed, the trained joint latent ranking embedding model can be applied to a test video clip for human action recognition. As a result, the relatedness scores corresponding to all the known and unseen action labels in a label collection are achieved by using both visual and semantic models (c.f. Section 3.4), as illustrated in the right box of Figure 1.

3.2. Long Short Term Memory

As described previously, the LSTM layer plays a crucial rule in tackling two technical issues arising from weakly annotated data for visual embedding in our proposed framework. Here, we briefly review the mechanism of an LSTM unit to facilitate understanding our framework.

As shown in Figure 2, an LSTM unit consists of a memory cell reserving the historic information at previous time steps. The output in time step t is determined by the current input x_t and the activation value in this memory cell. Three gates, *input*, *output* and *forgetting* gates, are used to control the information flow in the LSTM unit. The information flow in an LSTM unit [44, 46] is formulated as follows:

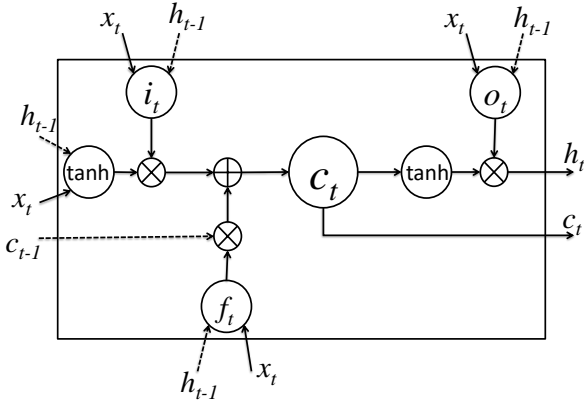


Figure 2: LSTM unit and its information flow.

$$\begin{aligned}
i_t &= \sigma(W_{xi}x_t + W_{hi}h_{t-1} + b_i), \\
f_t &= \sigma(W_{xf}x_t + W_{hf}h_{t-1} + b_f), \\
c_t &= f_t c_{t-1} + i_t \tanh(W_{xc}x_t + W_{hc}h_{t-1} + b_c), \\
o_t &= \sigma(W_{xo}x_t + W_{ho}h_{t-1} + b_o),
\end{aligned}$$

and

$$h_t = o_t \tanh(c_t),$$

where $\sigma(\cdot)$ and $\tanh(\cdot)$ are the sigmoid and the tangent hyperbolic functions, and i, f, o refer to input, forget and output gates, respectively. c is the cell state, representing the information stored in the cell, and h is output of the LSTM unit or input receiving from the output of another LSTM unit retroactively). W 's and b 's are weight matrices and bias vectors associated to different gates.

When LSTM units are used in a hidden layer as done in our framework, the dimension of hidden vector h is determined by the number of LSTM units used in this hidden layer. Hence, it is a hyper-parameter that has to be tuned on a given dataset. The dimensions of weight matrices W are thus determined by the dimensions of output vector h and input x . The gates i, f, o and the cell state c are of the same dimension as that of h .

3.3. Joint Latent Ranking Embedding Learning

Now we present the joint latent ranking embedding learning in our proposed framework. To facilitate our presentation, we summarize the notations used in this paper in Table 1.

3.3.1. General Description

Given a training set of weakly annotated video clips, $\mathcal{D} = \{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^{|\mathcal{D}|}$, where \mathbf{x}_i is the visual input and $\mathbf{y}_i \in \{+1, -1\}^{|\mathcal{C}^{Tr}|}$ is its binary target label vector in the i -th example: $+1/-1$ element indicates the presence/absence of a specific action belonging to \mathcal{C}^{Tr} in \mathbf{x}_i .

For a video instance \mathbf{x}_i in \mathcal{D} , we divide it evenly into T segments¹, segment-level visual representations are extracted,

Table 1: Nomenclature.

Notation	Description
$\mathcal{D}, \mathcal{D}_T$	training, test datasets
$ \cdot , \ \cdot\ _1$	cardinality of a set, L_1 norm of a vector
$\mathcal{C}, \mathcal{C}^{Tr}, \mathcal{C}^U$	label collection, training and unseen class label subsets
E^v, E^s	visual and semantic embedding space
\mathbf{x}_{it}	visual representation for the t -th segment of the i -th example
\mathbf{x}_i	collection of all the segment-level visual representations of the i -th example
\mathbf{s}_c	semantic representation for the c -th label
\mathbf{Y}	binary target label matrix of training dataset
\mathbf{y}_i	binary target label set of the i -th example, i.e., the i -th column of matrix \mathbf{Y}
\mathbf{y}^c	binary indicator vector of the c -th label appearing examples, i.e., the c -th row of matrix \mathbf{Y}
E^v, \mathbf{e}_i^v	visual embedding matrix and the column vector for i -th video clip
E^s, \mathbf{e}_c^s	semantic embedding matrix and the column vector for the c -th label
d_x, d_s, d_e	dimensions of visual, semantic, latent embedding space
\mathbf{o}_i	relatedness scores between the i -th example and all the candidate labels in visual model
\mathbf{o}^c	relatedness scores between the c -th label and all the training examples in semantic model
ϕ_v, Θ_v	visual embedding function and parameters
ϕ_s, Θ_s	semantic embedding function and parameters
ϕ	$\phi = \{\phi_v, \phi_s\}$, mapping function for multi-label zero-shot recognition
$\mathcal{C}(\hat{\mathbf{x}}), \mathcal{L}(\hat{\mathbf{x}})$	the ground-truth label set of test instance $\hat{\mathbf{x}}$, the ranking list of all the labels predicted for $\hat{\mathbf{x}}$ in terms of the relatedness scores
$\mathcal{D}^c, \mathcal{L}^c$	Set of test instances of which ground-truth label sets include the c -th label, the ranking list of all the test instances in terms of the relatedness scores for the c -th label

which are collectively denoted by $\{\mathbf{x}_{i1}, \mathbf{x}_{i2}, \dots, \mathbf{x}_{iT}\}$. At the t -th time step, the segment representation \mathbf{x}_{it} is fed into a hidden LSTM layer and processed by this LSTM layer and two subsequent fully-connected layers of linear activation functions (c.f. Figure 1). The latent visual embedding of the t -th segment, \mathbf{e}_{it}^v , is obtained as follows:

$$\mathbf{e}_{it}^v = \phi_v(\mathbf{x}_{it}; \Theta_v). \quad (1)$$

Here, ϕ_v is the visual embedding function implemented by the parametric visual model and Θ_v is a collective notation of all the parameters in this model, including weights and biases involved in this deep network.

Likewise, as depicted in Figure 1, the c -th label in a label collection is first represented by a specific semantic representation, \mathbf{s}_c , that is fed to the semantic embedding function, ϕ_s ,

¹A segment refers to a volume of multiple consecutive frames.

implemented by the parametric semantic model of which all the parameters are denoted by Θ_s collectively. Thus, the semantic embedding, \mathbf{e}_c^s , of the c -th label is

$$\mathbf{e}_c^s = \boldsymbol{\phi}_s(\mathbf{s}_c; \Theta_s). \quad (2)$$

A score layer is employed in each of the visual and the semantic models. In the visual model, the score layer takes the outputs of the visual embedding layer at all the time steps to yield the relatedness scores regarding all the labels for \mathbf{x}_i with a dot product between the visual embedding of each segment in \mathbf{x}_i and the semantic embedding of all the labels in a label collection:

$$\mathbf{o}_{it} = \langle \mathbf{e}_{it}^v, \mathbf{E}^s \rangle, \quad (3)$$

where $\mathbf{E}^s \in \mathbb{R}^{d_e \times |C^{Tr}|}$ is a collective notation of the semantic embedding of all the labels. Here, $\langle \mathbf{a}, \mathbf{B} \rangle = \mathbf{a}^T \mathbf{B}$ is a vectorial notation of the dot product between a vector, \mathbf{a} , and each column of a matrix, \mathbf{B} . Then the relatedness scores between this video instance and different labels are achieved by averaging over the scores on all the segments of this video instance:

$$\mathbf{o}_i = \frac{1}{T} \sum_{t=1}^T \mathbf{o}_{it}. \quad (4)$$

Likewise, the relatedness scores between different video instances and the c -th label in the label collection, $\mathbf{o}^c \in \mathbb{R}^{|\mathcal{D}| \times 1}$, is estimated in the same manner as done in the visual model based on the visual embedding of those video instances and the semantic embedding of the c -th label. Thus, the i -th element of \mathbf{o}^c , the relatedness score regarding the i -th visual instance is

$$o_i^c = \frac{1}{T} \sum_{t=1}^T \langle \mathbf{e}_{it}^v, \mathbf{e}_c^s \rangle. \quad (5)$$

For the joint latent ranking embedding learning, we need to optimize the parameters in the visual and the semantic models with training data and proper rank loss functions (the technical details are presented in Section 3.3.2). Assume that $l_v(\cdot, \cdot)$ and $l_s(\cdot, \cdot)$ are two loss functions with respect to the visual and the semantic model, the joint latent ranking embedding learning is boiled down to simultaneously solving the following optimization problems:

$$\Theta_v^* = \operatorname{argmin}_{\Theta_v} \sum_{i=1}^{|\mathcal{D}|} l_v(\mathbf{o}_i, \mathbf{y}_i), \quad (6)$$

$$\Theta_s^* = \operatorname{argmin}_{\Theta_s} \sum_{c \in C^{Tr}} l_s(\mathbf{o}^c, \mathbf{y}^c). \quad (7)$$

Here, the binary indicator vector $\mathbf{y}^c \in \mathbb{I}^{1 \times |\mathcal{D}|}$ is a row vector in the target label matrix $\mathbf{Y} \in \mathbb{I}^{|\mathcal{D}| \times |\mathcal{D}|}$ of a training dataset, \mathcal{D} , and elements of +1 in \mathbf{y}^c indicate that the c -th label appears in the target label sets of the corresponding training examples in \mathcal{D} . The binary indicator vector $\mathbf{y}_i \in \mathbb{I}^{|\mathcal{D}| \times 1}$ is a column vector in \mathbf{Y} , and elements of +1 in \mathbf{y}_i refers to those labels in the target label set associated with the i -th training example in \mathcal{D} . The value of elements corresponding to irrelevant visual input in \mathbf{y}^c or labels in \mathbf{y}_i is always set to -1.

3.3.2. Rank Loss Function

As described in Section 3.1, multi-label zero-shot learning needs to establish a mapping that outputs a label-relatedness score list for a video input where the scores of the relevant labels should be ranked higher than those of irrelevant ones. In previous studies, various rank loss functions have been developed for ranking-based learning [47]. To demonstrate the effectiveness of our joint latent ranking embedding framework, we adopt two simple yet typical rank loss functions, *RankNet loss* [48] and the margin-based *hinge rank loss* [49], in our work although other rank loss functions [47] may be employed in our framework as well. RankNet loss [48] provides a generic loss function for ranking-based learning from a probabilistic perspective, while hinge rank loss [49] was originally proposed for structural SVMs and has been widely used in different tasks including single-label zero-shot learning, e.g., [23, 18].

Nevertheless, we observe the following phenomenon in our experiments when using the original RankNet and hinge rank losses. By using only a ranking constraint in either of two rank losses, all the labels are considered independently and treated equally so that the less frequently used relevant labels might be overlooked during learning. Moreover, two rank losses generally make use of pairwise constraints to explore a relationship between labels associated with an instance explicitly. However, the relatedness scores in such rank losses are not bounded and could hence vary across different examples. Thus, some “difficult” pairs of labels are likely to incur a larger cost that predominates the overall loss, which could make the learning biased to those pairs of labels only. Furthermore, relatedness scores may vary in a large range for different training examples even though proper ranking relationships among them are established, which results in the poor performance. Motivated by the above observation, we introduce a regularization term to RankNet and hinge rank losses to overcome those problems.

For the target label set expressed with binary indicators, \mathbf{y}_i , in the i -th training example, $(\mathbf{x}_i, \mathbf{y}_i)$, the elements of +1 indicate all the labels relevant to \mathbf{x}_i while elements of -1 express all the remaining labels irrelevant to \mathbf{x}_i in terms of all the known actions in C^{Tr} . Likewise, \mathbf{y}^c , a binary indicator in $\{+1, -1\}$ regarding whether the c -th action appears in training examples in \mathcal{D} , can be handled in the same manner. Thus, the relatedness scores of \mathbf{x}_i to its positive and negative labels, \mathbf{o}_i , are achieved with Eqs. (2) and (3), and the relatedness scores of the c -th label to all the training examples, \mathbf{o}^c , are calculated with Eqs. (1) and (5). Based on the above quantities, we can define our regularized rank loss functions, $l_v(\mathbf{o}_i, \mathbf{y}_i)$ and $l_s(\mathbf{o}^c, \mathbf{y}^c)$.

Formally, we define the regularized RankNet loss function for visual embedding of \mathbf{x}_i as follows:

$$l_v(\mathbf{o}_i, \mathbf{y}_i) = \omega_i \left(\sum_{p \in C_i^{Tr+}} \sum_{q \in C_i^{Tr-}} \log(1 + \exp(o_{iq} - o_{ip})) + \sum_{j \in C^{Tr}} \log(1 + \exp(-y_{ij} o_{ij})) \right), \quad (8)$$

where $\omega_i = (|C_i^{Tr+}| \cdot |C_i^{Tr-}| + |C^{Tr}|)^{-1}$ normalizes this per-instance regularized rank loss. Corresponding to the elements of +1 and -1 in \mathbf{y}_i , C_i^{Tr+} and C_i^{Tr-} denote two subsets of relevant

and irrelevant labels to \mathbf{x}_i , respectively. Intuitively, minimizing the first term in Eq. (8) ensures that all the labels relevant to \mathbf{x}_i are ranked ahead of those irrelevant to \mathbf{x}_i . The second term in Eq. (8) plays a regularization role; minimizing this term during learning promotes the relatedness scores by enlarging the relatedness scores to the relevant labels as well as diminishing those to the irrelevant ones simultaneously, which tackles the problems observed in our experiments.

Likewise, we define the regularized RankNet loss function for semantic embedding of label c as follows:

$$l_s(\mathbf{o}^c, \mathbf{y}^c) = \omega_c \left(\sum_{p \in \mathcal{D}^{c+}} \sum_{q \in \mathcal{D}^{c-}} \log(1 + \exp(o_q^c - o_p^c)) + \sum_{j \in \mathcal{D}} (1 + \exp(-y_j^c o_j^c)) \right), \quad (9)$$

where $\omega_c = (|\mathcal{D}^{c+}| \cdot |\mathcal{D}^{c-}| + |\mathcal{D}|)^{-1}$ normalizes the per-class regularized rank loss. \mathcal{D}^{c+} and \mathcal{D}^{c-} are the positive and the negative training example subsets, respectively, regarding the c -th label. With the same treatment as used in Eq. (8), minimizing Eq.(9) ensures that the video instances conveying the action of the c -th label are ranked above all those without this action. Moreover, those video instances conveying the action of the c -th label, indicated by $y_j^c = +1$, tend to have as high relatedness scores as possible while all other video instances without this action, indicated by $y_j^c = -1$, tend to have as low relatedness scores as possible.

Similarly, we define a regularized hinge rank loss function for visual embedding of \mathbf{x}_i as follows:

$$l_v(\mathbf{o}_i, \mathbf{y}_i) = \omega_i \left(\sum_{p \in C_i^{Tr+}} \sum_{q \in C_i^{Tr-}} \max(0, m - o_{ip} + o_{iq}) + \sum_{j \in C^{Tr}} \max(0, m - y_{ij} o_{ij}) \right), \quad (10)$$

where $\omega_i = (|C_i^{Tr+}| \cdot |C_i^{Tr-}| + |C^{Tr}|)^{-1}$ and m is a pre-specified margin. Thus, minimizing the first term in Eq. (10) ensures that all the labels relevant to \mathbf{x}_i are ranked ahead of those irrelevant to \mathbf{x}_i with a pre-specified margin, m . The second term in Eq. (10) plays a regularization role; minimizing this term during learning promotes the margin-based relatedness scores by enlarging the relatedness scores to the relevant labels as well as diminishing those to the irrelevant ones simultaneously.

Likewise, we define a regularized hinge rank loss function for semantic embedding of label c as follows:

$$l_s(\mathbf{o}^c, \mathbf{y}^c) = \omega_c \left(\sum_{p \in \mathcal{D}^{c+}} \sum_{q \in \mathcal{D}^{c-}} \max(0, m - o_p^c + o_q^c) + \sum_{j \in \mathcal{D}} \max(0, m - y_j^c o_j^c) \right), \quad (11)$$

where $\omega_c = (|\mathcal{D}^{c+}| \cdot |\mathcal{D}^{c-}| + |\mathcal{D}|)^{-1}$ and m is a pre-specified margin. With the same treatment as used in Eq. (10), minimizing Eq.(11) ensures that the video instances conveying the action of the c -th label are ranked ahead of all those without this action with a margin, m . Moreover, those video instances conveying the action of the c -th label, indicated by $y_j^c = +1$, tend

to have as high relatedness scores as possible while all other video instances without this action, indicated by $y_j^c = -1$, tend to have as low relatedness scores as possible.

As a result, we can employ either our regularized RankNet loss functions in Eqs. (8) and (9) or the regularized hinge rank loss functions in Eqs. (10) and (11) to train visual and semantic embedding models in our framework.

3.3.3. Learning Algorithm

As formulated in Eqs. (6) and (7), learning is going to find the optimal parameters, Θ_v^* and Θ_s^* , by minimizing two loss functions, $l_v(\mathbf{o}_i, \mathbf{y}_i)$ and $l_s(\mathbf{o}^c, \mathbf{y}^c)$, defined in Section 3.3.2. However, the relatedness scores required in $l_v(\mathbf{o}_i, \mathbf{y}_i)$ regarding the visual model involve the output of the semantic model, E^s , and vice versa (c.f. Figure 1). Moreover, $l_v(\mathbf{o}_i, \mathbf{y}_i)$ requires the relatedness scores between all the candidate labels and each of training examples, while $l_s(\mathbf{o}^c, \mathbf{y}^c)$ needs the relatedness scores between all the training examples and each of all the action labels in C^{Tr} . Thus, our optimization problems are very complex and unsolvable simultaneously with commonly used local search methods, e.g., gradient-descent based methods.

Motivated by the works dealing with similar optimization problems, e.g., [50, 51], we come up with a learning algorithm to train the visual and the semantic models alternately during learning. In our alternate learning strategy, our learning algorithm begins with randomly initializing the parameters in the semantic model and then use the initialized parameter to generate the initial semantic embedding. By using the initial semantic embedding in $l_v(\mathbf{o}_i, \mathbf{y}_i)$, the visual model can be trained with a local search method such as the mini-batch stochastic gradient decent method. After one epoch, the current parameters in the visual model are frozen and used to generate the visual embedding for all the examples. By using the current visual embedding in $l_s(\mathbf{o}^c, \mathbf{y}^c)$, the semantic model is trained in the same manner. This alternate learning process carries on until a stopping condition is satisfied. The details of this *alternate* learning algorithm is described in Algorithm 1.

It is worth stating that two rank loss functions defined for visual and semantic model are related and the optimisation of one model would naturally promote the other towards its optimal solution. Hence, our alternate learning algorithm may converge after running finite epochs with the same properties held for similar methods [50, 51].

3.4. Multi-Label Zero-Shot Recognition

Once the joint latent ranking embedding learning is completed, we obtain a mapping function: $\phi(\mathbf{x}, \mathbf{c}) = \{\phi_v(\mathbf{x}|\Theta_v^*), \phi_s(\mathbf{c}|\Theta_s^*)\}$ where $\phi_v(\mathbf{x}|\Theta_v^*)$ and $\phi_s(\mathbf{c}|\Theta_s^*)$ are the visual and the semantic embedding functions implemented by the trained visual and semantic models, respectively. Then, we can use this mapping function for multi-label zero-shot human action recognition.

For recognition, we first extract the semantic representations of all the labels, including both *known* and *unseen* labels, in a considered label collection: $\mathbf{s}_c, \forall c \in C; C = C^{Tr} \cup C^U$ and $C^{Tr} \cap C^U = \emptyset$. By using the semantic embedding function, we achieve the semantic embedding of all the labels: $\hat{\mathbf{e}}_c^s =$

Algorithm 1 Joint Latent Ranking Embedding Learning

Input: Randomly initialize parameters, Θ_v^0 and Θ_s^0 , in the visual and the semantic models, respectively; extract the visual representations of training example, \mathbf{x}_i , $i = 1, \dots, N$, and the semantic representations of all the training labels, \mathbf{s}_c , $\forall c \in C^{Tr}$; input the target label matrix of the training set, Y ; pre-set the dimensionality of joint latent ranking embedding space, d_e .

Output: Optimal model parameters: Θ_v^* and Θ_s^* .

- 1: Generate the initial semantic embedding $\phi_s(\mathbf{s}_c; \Theta_s^0)$, $\forall c \in C^{Tr}$; $t \leftarrow 0$.
 - 2: **repeat**
 - 3: $t \leftarrow t + 1$;
 - 4: $\Theta_v^t = \arg\min_{\Theta_v} \sum_{i=1}^N l_v(\mathbf{o}_i, \mathbf{y}_i)$ with the current semantic embedding for one epoch;
 - 5: Generate the visual embedding with the current visual model, $\phi_v(\mathbf{x}_i; \Theta_v^t)$, $i = 1, \dots, N$;
 - 6: $\Theta_s^t = \arg\min_{\Theta_s} \sum_{c \in C^{Tr}} l_s(\mathbf{o}^c, \mathbf{y}^c)$ with the current visual embedding for one epoch;
 - 7: Generate the semantic embedding with the current semantic model $\phi_s(\mathbf{s}_c; \Theta_s^t)$, $\forall c \in C^{Tr}$;
 - 8: **until** Stopping condition is met.
 - 9: $\Theta_v^* \leftarrow \Theta_v^t$ and $\Theta_s^* \leftarrow \Theta_s^t$.
-

$\phi_s(\mathbf{s}_c | \Theta_s^*)$, $\forall c \in C$. For a test video clip, we divide it into T segments and extract its segment-level representations collectively denoted by $\hat{\mathbf{x}} = \{\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, \dots, \hat{\mathbf{x}}_T\}$. Technical details for extracting semantic and visual representations can be found in Section 4.2. By feeding $\hat{\mathbf{x}}$ to the visual embedding function, we achieve its visual embedding: $\hat{\mathbf{e}}^v = \phi_v(\hat{\mathbf{x}} | \Theta_v^*)$. Thus, the relatedness scores between this test video clip, $\hat{\mathbf{x}}$, and all the actions in the considered label collection C , including known and unseen labels during learning, is achieved by

$$S(\hat{\mathbf{x}}, c) = \langle \hat{\mathbf{e}}^v, \hat{\mathbf{e}}_c^s \rangle, \quad \forall c \in C. \quad (12)$$

Finally, we achieve a ranking action label list, $L(\hat{\mathbf{x}})$, for this test video clip by sorting its relatedness scores measured against all the labels in C with Eq. (12):

$$L(\hat{\mathbf{x}}) = \{c_i\}_{i=1}^{|C|}, \quad (13)$$

where $\forall c_i, c_j \in C$, $\text{Score}(\hat{\mathbf{x}}, c_i) \geq \text{Score}(\hat{\mathbf{x}}, c_j)$ if $i < j$.

In our experiments regarding the use of two different rank losses in our framework, we observe that the regularized RankNet and hinge rank losses often behave differently in several evaluation scenarios (c.f. Section 4). Although two rank losses generally yield the comparable performance overall, a closer look suggests that those correctly recognized video instances are quite different when two different rank losses are used in our framework, respectively. Hence, we would employ a simple fusion method to exploit the complementary aspect resulted from the use of two different rank losses. In order to fuse the results yielded by the models trained with two different rank losses, we first normalize the relatedness scores of $\hat{\mathbf{x}}$ to the considered label collection, C , generated by each of two models as

follows:

$$\tilde{S}(\hat{\mathbf{x}}, c) = \frac{S(\hat{\mathbf{x}}, c) - S_{\min}}{S_{\max} - S_{\min}}, \quad (14)$$

where S_{\max} and S_{\min} are the highest and lowest related scores of video clips, respectively, measured on a test set. Let $\tilde{S}^{(Reg)}(\hat{\mathbf{x}}, c)$ and $\tilde{S}^{(Hinge)}(\hat{\mathbf{x}}, c)$ denote the normalized relatedness scores yielded by two models trained with our regularized rank and the hinge rank losses, respectively. Then, the fused relatedness scores, $\tilde{S}^{(Fusion)}(\hat{\mathbf{x}}, c)$ is simply an average between $\tilde{S}^{(RankNet)}(\hat{\mathbf{x}}, c)$ and $\tilde{S}^{(Hinge)}(\hat{\mathbf{x}}, c)$; i.e.,

$$\tilde{S}^{(Fusion)}(\hat{\mathbf{x}}, c) = \frac{\tilde{S}^{RankNet}(\hat{\mathbf{x}}, c) + \tilde{S}^{Hinge}(\hat{\mathbf{x}}, c)}{2}. \quad (15)$$

Based on the fused relatedness scores, a ranking action label list, $L^{(Fusion)}(\hat{\mathbf{x}})$, is achieved in the same manner as specified in Eq.(13) for any test video clip, $\hat{\mathbf{x}}$.

4. Experimental Setting

In this section, we describe our experimental design and settings, including datasets, visual and semantic representations, model learning, evaluation scenarios and criteria used in our experiments. Moreover, we design a number of comparative experiments to exhibit the gain resulting from different components in our framework and to demonstrate the effectiveness of our framework by a comparison to several state-of-the-art multi-label ZSL methods that could be applied to general human action recognition.

4.1. Datasets and Splits

We first describe datasets and their split settings used in our experiments for simulation of multi-label ZSL scenarios.

4.1.1. Datasets

To evaluate our framework, we employ two publicly available video datasets: *Breakfast* [52] and *Charades* [4], in our experiments. In both datasets, at least two actions are involved in each video clip and the duration of each video clip is relatively long, which implies the temporal coherence information may be explored and exploited in human action recognition. Hence, both datasets are suitable to evaluate weakly annotated multi-label human action recognition. Below, we summarize the main aspects of two video datasets.

Breakfast: In this dataset [52], there are 1,989 video clips totally, where a video clip conveys several cooking actions. Totally, there are 49 cooking actions (excluding the “silence” label), such as “stirring”, “pouring_milk” and “opening_the_fridge”. Those actions are performed by 52 people in different kitchens. Although this dataset is not collected especially for multi-label human action recognition, we would use it as a proof-of-concept test bed.

Charades: This dataset [4] is collected from hundreds of people recording videos in their own home especially for video-based human activity analysis in daily lives. Hence, it is very challenging for multi-label human action recognition. In this

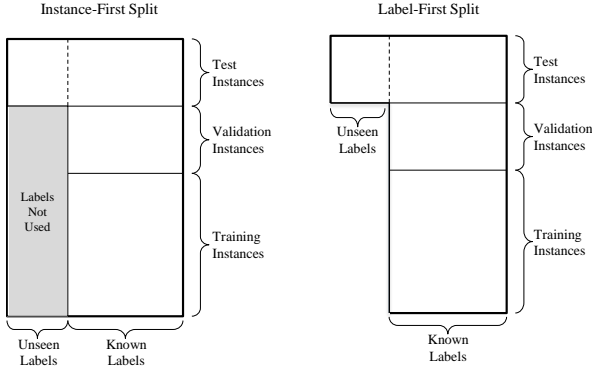


Figure 3: Two different data split settings used in our experiments. In the *instance-first* split (left plot), a data set is simply split into three mutually exclusive subsets: training, validation and test subsets. Any unseen labels associated with instances in training and validation subsets (shaded area) are removed from their target label sets before being used in learning. In the *label-first* split (right plot), a number of labels are first specified as unseen labels. Instances associated with any of unseen labels form a test subset. The remaining data of known labels are further divided into training and validation subsets to be used in learning.

dataset, there are 9,848 video clips involving 157 different human actions totally, acting out casual everyday activities. An average duration of video clips is around 30 seconds and an average number of actions involved in a video clip is 6.8. Those actions are performed by 267 people from three continents, and more than one person appear in over 15% of all the video clips. The raw video data (scaled to 480p) are used in our experiments, which are available from the Charades project page².

4.1.2. Dataset Splits

To simulate a zero-shot scenario, we need to split a dataset into training and test sets where a training set contains examples associated with only known classes while a test set has test instances involving at least one unseen class. Unlike single-label ZSL where a dataset is automatically split into training and test sets once unseen classes are specified, the dataset split issue in multi-label ZSL becomes much more complicated. In our experiments, we make two different split settings, *instance-first split* (IFS) and *label-first split* (LFS), as illustrated in Figure 3.

Instance-First Split

This is a commonly used data split setting in all the existing multi-label ZSL works, e.g., [22, 32, 6]. In this setting, instances in a dataset is first split into training, validation and test subsets. The training and the validation subsets are used for parameter estimation and hyper-parameter tuning, the dimension of latent embedding space d_e and the number of iterations of Algorithm 1 for our model. The test set that may or may not involve unseen labels is reserved for performance evaluation. Then, we divide the action label collection into mutually exclusive known and unseen label sets. Before learning, any unseen labels in the target label set of an instance in the training and the validation subsets are removed as shown in the left plot of

Figure 3. In other words, only known labels in the target label sets of an instance in those two datasets are used in learning. It is worth clarifying that unlike single-label ZSL, it is often infeasible to simulate a ZSL scenario by manipulating the validation set due to insufficient data in two datasets used in our experiments. Hence, the validation for hyper-parameter tuning in our experiments has to follow the typical protocol used in multi-label learning [3].

To split the Breakfast dataset with this setting, we adopt the pre-split by data collectors [52], where the video clips of 13 people are reserved for test. We further divide the rest video clips for training and validation: video clips of 32 people for training and the remaining video clips of seven people for validation. As a result, the numbers of video clips for training, validation and test are 1,196, 126 and 667, respectively. Then we randomly split the 49 labels into known and unseen labels: 10 labels reserved as unseen labels and the rest 39 as known labels.

For the Charades dataset, we also adopt its pre-split provided by data collectors [4], where 7,985 and 1,863 video clips are used for training and test, respectively. We further divide training data into two subsets: 6,385 for training and 1,600 for validation in our experiments. Then we randomly choose 40 out of 157 human actions as unseen classes and the rest 117 human actions are hence known actions.

Label-First Split

Although the instance-first data split setting is widely used in multi-label ZSL, it suffers from a fundamental limitation. It is well known that multiple labels together could frame a specific concept and removing any label from this label cohort may lead to a less accurate semantic meaning and biases in learning. Furthermore, the instance-first split allows for accessing to visual features of instances involved in unseen actions. To overcome this limitation, we propose a novel data split setting for multi-label ZSL named *label-first* split. In this new setting, all the labels in a label collection used in a dataset is first divided into two mutually exclusive subsets: known and unseen labels. Then, all the instances having any unseen labels are reserved for test and the rest instances of known labels only are further divided into two subsets for training and validation, as shown in the right plot of Figure 3. Due to sparsity of training data, the validation in the label-first split also adopts the protocol used in multi-label learning [3].

To split the Breakfast dataset with this setting, we randomly choose 10 labels for unseen labels and the rest 39 labels are designated as known labels accordingly. Hence, this dataset is naturally split into two sets for training and test. The training data are further divided for training and validation. For validation, we randomly choose 200 instances from the training data. Likewise, the Charades dataset is split by using 20 randomly chosen label for unseen labels. Thus, the remaining 137 labels become known labels. From the instances of known labels, we randomly choose 1,000 instance used for validation. It is worth stating that the current datasets do not allows for reserving a large number of classes as unseen classes in either the IFS or the LFS setting. In the IFS, the more labels reserved as un-

²<http://allenai.org/plato/charades/>

Table 2: Information on two different data split settings.

Dataset	Split Method	# Training Inst.	# Validation Inst.	# Test Inst.	# Known Labels	# Unseen Labels
Breakfast	Instance-first	1196	126	667	39	10
	Label-first	1019/917/823	200	770/872/966	39	10
Charades	Instance-first	6385	1600	1863	117	40
	Label-first	4580/4176/3987	1000	4268/4672/4861	137	20

seen labels, the less accurate mapping learned from visual to semantic domains due to the existence of visual features of unseen actions and a lack of the corresponding action labels in such training examples. In the LFS, the more labels reserved as unseen labels, the fewer training examples available. Hence, the training examples do not convey the essential information required in learning.

For reliability, we repeat our experiments on each dataset under each split setting for three trials. During a trial, training data given in the pre-split of each dataset is randomly divided into training and validation subsets with the instance-first split setting, and a known/unseen label split on each dataset is chosen randomly with the label-first split setting. For clarity, we summarize the data split information³ on two datasets in Table 2.

4.2. Visual and Semantic Representations

In our experiments, we use visual representations extracted with the existing C3D deep network [53] and word vectors as semantic representations [39].

As suggested by [53], the C3D features are extracted for a segment of 16 frames with eight frames overlapping between two adjacent segments. Thus, a training/test video clip is always divided into T segments with the treatment as follows. To ensure that each video clip can be divided into T segments, any video clip must have $8 * (T + 1)$ frames. To this end, we simply down-sample those video clips of more frames with a proper sampling rate so that T C3D feature vectors can be extracted and collectively form a *segment-based* visual representation for this video clip. When a video clip has fewer frames, we first extract C3D feature vectors from those frames in this video clip and then pad all-zero vectors to the visual representation until there are T feature vectors. Also, we can convert T feature vectors into a holistic *instance-level* visual representation via averaging those T feature vectors. By using such an instance-level visual representation in our comparative study, we would demonstrate a performance gain benefiting from exploring/exploiting temporal coherence information underlying segments in a video clip. In our experiments, the segment-based visual representation is always used in our model while the instance-level visual representation is used in the baseline and the state-of-the-art models (c.f. Section 4.6) unless a different setting is specified. Based on a cross-validation experiment, we choose $T = 300$ for Breakfast and $T = 20$ for Charades. Although only C3D features are used in our experiments, it is worth mentioning that other kinds of visual representations,

e.g., IDT features [54] and deep image features extracted on a frame basis, can also be used straightforwardly in our framework.

In our experiments, we adopt Word2Vec as our semantic representation. Word2Vec was trained with a skip-gram neural network on the Google News dataset of 100 billion words [39]. As a result, one action label is represented by a 300-dimensional word vector. Although only 300-dimensional word vectors are used in our experiments, word vectors of different dimensionality may be used, and moreover, other semantic representations, e.g., attributes, can be used in our framework without any difficulty if available.

It is worth emphasizing that the same treatment described above is applied in both the learning and the recognition phases to extract visual and semantic representations.

4.3. Model Learning

In our experiments, model learning is implemented on Keras [55], a high-level neural networks library, running on top of either TensorFlow or Theano. As we use two neural networks to carry out the visual and the semantic models (c.f. the left box in Figure 1), we need to decide the specific network architectures and relevant hyper-parameters on two datasets during the model learning. The optimal hyper-parameters are found by a grid-based search via a cross-validation procedure. The Adam [56], a stochastic optimisation method, is used for training our model with its default configuration.

The visual model takes a sequence of segment-level C3D representations of $d_x = 4,096$ features as input to the LSTM layer where there are N_1 LSTM units. To improve the generalization, we also apply the *dropout* procedure [57] to the LSTM layer where a dropout rate needs specifying. The output of LSTM units are fed to a fully connected dense layer of N_2 neurons, and the output of this dense layer are further fed to the visual embedding layer of d_e neurons. During learning, there are no hyper-parameters involved in the score and the average pooling layer in the visual model.

As described in Section 3.1, the semantic model is carried out by a fully-connected three-layer feed-forward neural network. The word vectors of $d_s = 300$ dimensions are first input to a hidden layer of N_1 *ReLU* units. Subsequently, the output of this hidden layer are fed to the semantic embedding layer of d_e linear units. Note that for joint latent ranking embedding learning, the dimension of the semantic embedding space is set to the same of the visual embedding space in our experiments. Likewise, there are no hyper-parameters involved in the score and the average pooling layer in the semantic model during learning.

³All the data splits and source code used in our experiments are available on our project website: <http://staff.cs.manchester.ac.uk/~kechen/MLZSHAR>.

4.4. Evaluation Scenarios

Multi-label zero-shot recognition is complex given the fact that a test instance may be associated with a label set including both known and unseen class labels. Thus, there are different evaluation scenarios in previous works [31, 6]; each focuses on a specific aspect. Following their settings [31, 6], we evaluate our framework along with other learning models used in our comparative study described later on in this section in three different scenarios:

Known-action only: In this setting, the performance is evaluated regarding only known (training) actions. This scenario boils down to the conventional supervised multi-label learning. In this circumstance, we no longer take any unseen action labels into account during test; for a test instance, its relatedness score ranking list contains only those regarding known action labels in C^{Tr} and any unseen action label in C^U in its ground-truth label set, if there is, will be removed such that the modified ground-truth set includes only known action labels in C^{Tr} .

Unseen-action only: In this setting, the performance is evaluated regarding only unseen (test) actions. This scenario boils down to a standard ZSL setting. In this situation, we no longer consider any known action labels; for a test instance, its relatedness score ranking list contains only those regarding unseen action labels in C^U and any known action label in C^{Tr} in its ground-truth label set, if there is, will be removed such that the modified ground-truth set includes only unseen action labels in C^U .

Generalized ZSL: In this setting, the performance is evaluated regarding all the actions of which labels appearing in a label collection C without considering if an action label is known or unseen during learning. This scenario has been named *generalized ZSL* in the machine learning community. In this situation, both known and unseen action labels are treated equally; for a test instance, its relatedness score ranking list contains those regarding all the action labels in C and the evaluation is made against its ground-truth label set that could be a mixture of known and unseen labels. It is worth highlighting that the generalized ZSL setting is required by multi-label zero-shot human action recognition in a real application.

4.5. Evaluation Metrics

There are a variety of evaluation metrics for multi-label learning. Depending on the output of a multi-label learning system, the evaluation metrics are generally divided into two types: *ranking-based* and *bipartition-based* metrics [32, 58]. Ranking-based metrics work for the situation that a learning system yields a ranking list of continuous-valued relatedness scores on all the candidate labels. In contrast, bipartition-based metrics are used when a learning system produces only a binary indicator vector for all the candidate labels, where 1/0 element expresses the presence/absence. Since our model yields a ranking list of continuous-valued relatedness scores, we employ two commonly used ranking-based metrics for performance evaluation [32, 58, 22, 6, 59], Instance-centric Mean Average Precision (I-MAP) and Label-centric Mean Average Precision

(L-MAP). In addition, we employ other metrics, *precision*, *recall* and F_1 score, which have also been used in the performance evaluation of multi-label learning [3, 60].

To facilitate our presentation, we first define the *precision-at-k* [61] in a generic form:

$$P@k(A, B) = \frac{1}{k} |A \cap B[1, \dots, k]|, \quad (16)$$

where A is a ground-truth set, B is a set of all the retrieved entities ranked in terms of relevance, and $B[1, \dots, k]$ indicates top k entities in B . Given a test dataset, $\mathbb{D}_T = \{\hat{\mathbf{x}}_i\}_{i=1}^{|\mathbb{D}_T|}$, a learning model yields a *label-based* ranking list for a test instance, $\hat{\mathbf{x}}_i \in \mathbb{D}_T$, in terms of its relatedness scores to all the labels in C (c.f. Eqs. (12) and (13)): $L(\hat{\mathbf{x}}_i) = \{c_j\}_{j=1}^{|C|}$, where $\forall c_p, c_q \in C$, $\text{Score}(\hat{\mathbf{x}}_i, c_p) \geq \text{Score}(\hat{\mathbf{x}}_i, c_q)$ if $p < q$. Let $C(\hat{\mathbf{x}}_i)$ denote the ground-truth label set of $\hat{\mathbf{x}}_i$. I-MAP over a test dataset \mathbb{D}_T is defined by

$$\text{I-MAP} = \frac{1}{|\mathbb{D}_T|} \sum_{i=1}^{|\mathbb{D}_T|} \frac{\sum_{c=1}^{|C|} P@c(C(\hat{\mathbf{x}}_i), L(\hat{\mathbf{x}}_i)) \delta(c, C(\hat{\mathbf{x}}_i))}{|C(\hat{\mathbf{x}}_i)|}, \quad (17)$$

where $\delta(c, C(\hat{\mathbf{x}}_i)) = 1$ if $c \in C(\hat{\mathbf{x}}_i)$ and $\delta(c, C(\hat{\mathbf{x}}_i)) = 0$ otherwise.

While I-MAP measures the accuracy in terms of test instances, L-MAP is used to evaluate the performance from a different perspective in light of candidate labels. Given a specific label $c \in C$, a model predicts the relatedness scores against the action specified by the c -th label for all the test instances in \mathbb{D}_T . Hence, we can achieve an *instance-based* ranking list for the c -th label, $L^c = \{\hat{\mathbf{x}}_{ij}\}_{j=1}^{|\mathbb{D}_T|}$, in terms of their relatedness scores against the c -th label where $\forall \hat{\mathbf{x}}_{ip}, \hat{\mathbf{x}}_{iq} \in \mathbb{D}_T$, $\text{Score}(\hat{\mathbf{x}}_{ip}, c) \geq \text{Score}(\hat{\mathbf{x}}_{iq}, c)$ if $p < q$. Let D^c denote the collection of those test instances of which their ground-truth label sets indeed include the c -th label. Thus, the L-MAP over a test dataset, \mathbb{D}_T , is defined by

$$\text{L-MAP} = \frac{1}{|C|} \sum_{c=1}^{|C|} \frac{\sum_{i=1}^{|\mathbb{D}_T|} P@i(D^c, L^c) \delta(\hat{\mathbf{x}}_i, D^c)}{|D^c|}, \quad (18)$$

where $\delta(\hat{\mathbf{x}}_i, D^c) = 1$ if $\hat{\mathbf{x}}_i \in D^c$ and $\delta(\hat{\mathbf{x}}_i, D^c) = 0$ otherwise.

Those widely used evaluation metrics in information retrieval have also been used in evaluating multi-label learning systems, e.g., [3, 60]. In our experiments, we adopt overall top- k precision, recall and F_1 score measured over a test dataset, \mathbb{D}_T , which are defined as follows:

$$\text{precision}(k) = \frac{\sum_{i=1}^{|\mathbb{D}_T|} P@k(C(\hat{\mathbf{x}}_i), L(\hat{\mathbf{x}}_i))}{k * |\mathbb{D}_T|}, \quad (19)$$

$$\text{recall}(k) = \frac{\sum_{i=1}^{|\mathbb{D}_T|} P@k(C(\hat{\mathbf{x}}_i), L(\hat{\mathbf{x}}_i))}{\sum_{i=1}^{|\mathbb{D}_T|} |C(\hat{\mathbf{x}}_i)|}, \quad (20)$$

$$F_1(k) = \frac{2 * \text{precision}(k) * \text{recall}(k)}{\text{precision}(k) + \text{recall}(k)}. \quad (21)$$

4.6. Comparative Study

In our experiments, we systematically conduct a comparative study from two different perspectives: ablation study and state-of-the-art models. As a result, a number of baseline systems are designed to demonstrate roles played by the main components in our framework while several state-of-the-art multi-label ZSL algorithms are adapted for human action recognition. For the comparative study, we evaluate each of different models on three evaluation scenarios with evaluation metrics described in Sections 4.4 and 4.5 under the exactly same conditions, including visual and semantic representations. As there are alternative pooling strategies that could be used to implement our framework, we further investigate those pooling strategies by comparing them to the average pooling used in our framework.

4.6.1. Baseline Models

To investigate the roles played by different component mechanisms employed in our framework, we design four baseline models, *random guess of scores*, *non-recurrent connection, without semantic embedding* and *randomized label representation*, by manipulating our framework with different purposes described as follows:

Random guess of scores (RGS): This is a general baseline that provides a lowest performance bound used for a reference to improvement made by a learning model. In our work, we randomly generate relatedness scores of all the candidate labels for a test instance. Then the performance of this baseline model is evaluated based on the random guess of scores. For reliability, we repeat the RGS process 100 times in our experiments and the statistics of the RGS performance including mean and standard error of mean (SEM) are reported.

Non-recurrent connection (NRC): In our framework, a LSTM layer of recurrent connections is employed to capture temporal coherence underlying sequential video data in the visual embedding learning. To examine the role played by the LSTM layer, we replace the recurrent connected layer with a fully connected layer without recurrent connections and keep all other components in our framework unchanged. By this setting, our model is converted into a baseline model named *non-recurrent connection*. During learning, obviously, this baseline model no longer explicitly makes use of the temporal dependency information underlying sequential segments in a video clip. Algorithm 1 is used directly for parameter estimation.

Without semantic embedding (WSE): In our framework, there is a semantic model for semantic embedding with the motivation that the use of a joint latent ranking embedding space narrows the semantic gap between visual and semantic domains and the zero-shot recognition should be done in the joint latent ranking embedding space. However, some existing works, e.g., Fast0Tag [6], do not learn a semantic embedding and the zero-shot recognition takes place directly in the semantic space. To examine the effectiveness of our semantic embedding, we come up with a baseline model named *without semantic embedding* by removing the semantic model from our framework. Thus, the original semantic

representations are used to replace the semantic embedding representations, \mathbf{E}^s , required by the score layer in the visual model, which is amount to mapping the visual space directly onto the original semantic space. As this baseline model has only the visual model, the learning becomes simpler; i.e., solving the optimization problem formulated in Eq. (6) based on the original semantic representation with the Adam [56]. It is worth clarifying that this baseline model is similar to Fast0Tag [6] apart from an LSTM-based visual embedding model and the segment-level visual representation used in this baseline model while a feed-forward neural network and instance-level visual representation are employed by Fast0Tag for visual embedding.

Randomized label representation (RLR): One of the most important issues in ZSL is exploring/exploiting the side information conveyed in the semantic domain. As our framework works for multi-label zero-shot recognition, we would investigate whether the semantic relatedness information encoded in the semantic embedding, inherited from the original semantic representations, is effectively used in knowledge transfer. To this end, we design another baseline model named *randomized label representation* by replacing the word vector of a label with a vector of the same dimensionality that is generated randomly and normalized with the l_2 norm to ensure that it has the same range as that of the word vector. Apparently, the semantic relatedness information no longer exists in such randomized label representations. For parameter estimation, Algorithm 1 is used directly via replacing the semantic representations of labels with the randomized label representations in training data.

4.6.2. State-of-the-Art Methods

Although, to the best of our knowledge, there exists no work in multi-label zero-shot human action recognition, we notice that there are a few multi-label ZSL algorithms. In our comparative study, we adopt and extend those multi-label ZSL algorithms for human action recognition for a thorough evaluation of our proposed framework. Below, we briefly describe those multi-label ZSL algorithms used in our experiments.

Direct Semantic Prediction (DSP): DSP is a well-known baseline model used in previous works for multi-label ZSL, e.g., [31]. DSP is derived from *direct attribute prediction* originally proposed for single-label ZSL [16]. The idea behind DSP is learning a mapping function $\phi : X \rightarrow \mathcal{S}$ from visual to semantic space directly for ZSL. By using the composition property of word vectors, given a multi-labelled video clip, we use the mean word vector achieved by averaging those word vectors of the labels associated with this video clip to be its semantic representation. Thus, the multi-label ZSL problem is boiled down to single-label ZSL. In our experiments, we employ support vector regressor models to learn the mapping function $\phi(\cdot)$. Given a test instance \mathbf{x} , the learned $\phi(\cdot)$ is used to predict its compositional semantic representation $\mathbf{f} = \phi(\mathbf{x})$, then the prediction scores regarding different labels are estimated by measuring distances between the predictions and the word vectors of all the labels in a label collection C : $\text{Score}(\mathbf{x}, c) = \langle \mathbf{f}, \mathbf{s}_c \rangle$, $\forall c \in C$,

which leads to a label-based ranking list, $L(\hat{\mathbf{x}})$, in terms of semantic relatedness.

Convex combination of Semantic Embedding (ConSE): ConSE is a ZSL algorithm proposed by [62], which can be naturally applied to multi-label ZSL. As same as formulated in DSP, ConSE also learns a mapping to predict a compositional semantic representation from the visual representation of a given video clip. Instead of learning a direct mapping function as DSP does, however, ConSE learns the conditional probabilities, $P(c|\mathbf{x})$ for $\forall c \in C^{Tr}$, regarding all the known actions with training data. For recognition, the compositional or collective semantic representation of a test instance $\hat{\mathbf{x}}$ is estimated by a convex combination of the semantic representations of top-5 known actions of the highest conditional probabilities. The combination weights are the l_2 normalized conditional probabilities of those top-5 known actions: $\hat{\mathbf{s}} = \sum_{c \in C^{Tr}} P(c|\hat{\mathbf{x}}) \mathbf{s}_c$. Thus, the prediction scores regarding all the labels in an action label collection C are $\text{Score}(\hat{\mathbf{x}}, c) = \langle \hat{\mathbf{s}}, \mathbf{s}_c \rangle$, $\forall c \in C$, which leads to a label-based ranking list, $L(\hat{\mathbf{x}})$.

COSTA: COSTA is a method proposed by [22] especially for multi-label zero-shot classification. In this method, multi-label classification is converted into a number of binary classification problems via a one-vs-rest setting. $|C^{Tr}|$ linear binary SVMs are trained based on the examples regarding $|C^{Tr}|$ known actions. Then the parameters of the SVM for an unseen label $c \in C^U$ is estimated by a weighted combination of the parameters of $|C^{Tr}|$ trained SVMs corresponding to known actions: $\mathbf{w}_c = \sum_{k=1}^{|C^{Tr}|} \alpha_k \beta_{ck} \mathbf{w}_k$. Here, \mathbf{w}_k is the parameters of the SVM regarding the k -th label in C^{Tr} and α_k is a combination coefficient regarding the importance of this SVM achieved via learning. $\beta_{ck} = \exp(-d_{ck}) / \sum_{j=1}^{|C^{Tr}|} \exp(-d_{cj})$, which is a factor indicating the relatedness between a unseen label, $c \in C^U$, and a known label, $k \in C^{Tr}$, and d_{ck} is the semantic distance between labels c and k measured via their word vectors. Thus, the SVM with the parameters \mathbf{w}_c can be used to predict the unseen label c for a given test instance. Note that our experimental results not reported in this paper due to the limited space suggest that learning α_k is not only time consuming but also yields the poorer performance than that where all $|C^{Tr}|$ SVMs are treated equally; i.e., $\alpha_k = 1$ for $k = 1, \dots, |C^{Tr}|$. Later on, we only report the best performance under this setting.

Fast0Tag: Fast0Tag is one of the latest state-of-the-art methods proposed for multi-label image tagging and multi-label ZSL [6]. The main idea behind Fast0Tag is learning a mapping function $\phi : X \rightarrow S$ from visual to semantic space for multi-label zero-shot tagging and recognition. Unlike DSP, a ranking-based loss function, *RankNet*, is used to train a deep network to carry out $\phi(\cdot)$ so that for an video clip, its relevant labels should be ranked ahead of those irrelevant ones. For recognition, the mapping function yields the predicted semantic representation, $\phi(\hat{\mathbf{x}})$, for a test instance, $\hat{\mathbf{x}}$. Then, we can achieve the relatedness scores to all the labels in an action label collection and the label-based ranking list as same as done in DSP and ConSE. In our experiments, we strictly follow the same settings suggested by [6].

Fast0Tag+: Our work presented in this paper suggests that

the use of learned semantic embedding leads to better performance than the use of the original semantic representations directly. To further investigate this idea, we make an extension of Fast0Tag by incorporating our semantic model into the Fast0Tag model and name our extension *Fast0Tag+*. As a result, Fast0Tag+ has an architecture resembling ours (c.f. the left box of Fig.1), where the visual model is carried out by the original Fast0Tag architecture while the semantic model is the same as ours presented in Section 3. The original rank loss functions in Fast0Tag are used and our alternate learning algorithm described in Algorithm 1 is used for parameter estimation. For recognition, the same procedure presented in Section 3.4 is used for a given test instance. Here, we argue that this extension would provide further evidence in examining the effectiveness of semantic embedding learning.

4.6.3. Pooling Strategy

To investigate the effect of different pooling strategies over temporal relatedness scores, we conduct a comparative experiment by replacing the average pooling with either the maximum pooling or the local average global maximum pooling in our framework. For the maximum pooling, Eq.(4) for the average pooling is thus altered to

$$\mathbf{o}_i = \max_{t=1}^T \mathbf{o}_{it}. \quad (22)$$

For the local average global maximum pooling, we firstly divide the T segments into T_s groups with a 50% overlap between two consecutive groups. As a result, there are $N_g = 2 * T / T_s$ consecutive segments in each group. We calculate the average score in each group and find the maximum as follows:

$$\mathbf{o}_i = \max_{t_s=1}^{T_s} \frac{1}{N_g} \sum_{t=(t_s-1)N_g/2+1}^{(t_s+1)N_g/2} \mathbf{o}_{it}. \quad (23)$$

Note that the local average global maximum pooling strategy is generic, and the average pooling and maximum pooling can be viewed as its special cases without between-group overlapping: the average pooling when $T_s = 1, N_g = T$ and the maximum pooling when $T_s = T, N_g = 1$, respectively. To make a thorough investigation, we set $T_s = 10, N_g = 60$ and $T_s = 20, N_g = 30$ as two experimental settings for Breakfast dataset. For Charades dataset, we set $T_s = 5, N_g = 8$ and $T_s = 10, N_g = 4$. All other experimental settings are kept the same for a fair comparison.

In our comparative study, the optimal hyper-parameters involved in baseline and state-of-the-art learning models are sought during their learning with the same cross-validation procedure as described in Section 4.3. Moreover, five state-of-the-art methods described above and ours are extensible to multi-label recognition straightforwardly; i.e., all the actions are known in advance and their training examples are available during learning. Thus, we also report the multi-label recognition performance, which not only extends our comparative study in a wider scope but also provides a benchmark to see how much the performance of each method is degraded in a zero-shot circumstance. For experiments in comparison of different pooling strategies, all the components and setting are kept unchanged except the pooling operations.

5. Experimental Results

In this section, we report the detailed experimental results in different settings and exemplify some typical test instances via visual inspection.

5.1. Results on Learning

We first report the experimental results regarding learning including optimal hyper-parameters for all the models used in our experiments and the evolution of the learning process for our model trained with our proposed alternate learning algorithm (Algorithm 1) under different data split settings.

As described in Section 4.3, we employ a grid-based search procedure via cross-validation to find out the optimal hyper-parameters in terms of both the loss used to train a model and the I-MAP performance (c.f. Section 4.5) as this metric directly evaluate the relatedness of a video instance to all the labels in a considered action label collection. We seek an optimal value from a set of candidate hyper-parameters involved in all different learning models used in our experiments with the exactly same procedure as follows:

Network architecture: The optimal architecture of neural networks in a model used in our experiments is investigated by tuning different number of neurons in each hidden layer. In our proposed model, there are totally four structural hyper-parameters. The number of hidden units in the LSTM layer is selected from the candidate set, $N_1 = 256, 512, 1024$. In the visual model, the number of neurons in the hidden layer above the LSTM layer is investigated with $N_2 = 1024, 2048$. In the semantic model, the number of neurons in the first hidden layer is selected from $N_1 = 300, 500, 700$. As a critical hyper-parameter in our algorithm, the dimension of latent embedding space d_e , the number of neurons in the visual/semantic embedding layers, is investigated by setting the candidate values, $d_e = 200, 500, 800$. For the *non-recurrent* baseline model, the number of neurons in the first hidden layer replacing the LSTM layer is chosen from $N_1 = 1024, 2048, 4096$. For Fast0Tag and Fast0Tag+, the number of neurons in the first and second hidden layers are selected from $N_1 = 2048, 4096, 8092$ and $N_2 = 1024, 2048$, respectively.

Learning rate: For all the neural networks in the proposed model, the baseline and the state-of-the-art models, candidate learning rates are $\{1e-2, 1e-4\}$ and $\{1e-4, 1e-6\}$ for the visual and the semantic models, respectively.

Number of epochs: Learning is stopped when the I-MAP performance on a validation set is no longer improved within the last 10 epochs and the loss reaches a low level on both training and validation sets. Then, the optimal model chosen is the one that yields the highest value of I-MAP on the validation set.

Dropout rate: The dropout rate used in the first layer of a neural model during learning is selected from $\{0, 0.5\}$.

Margin: The margin used in the hinge rank loss is selected for $m = 0.1, 1, 10$.

SVM hyper-parameters: In our comparative study, ConSE [62] and COSTA [22] employ a linear SVM for classification

and DSP [16] uses a linear SVR for regression. In our experiments, an optimal soft-margin value is sought from $C = 0.01, 1, 100$. For SVR, the percentage of support vectors is always set to $\varepsilon = 0.1$ as suggested in literature.

As a result, the resultant optimal hyper-parameter values in different experimental settings are summarized in Table 3.

To train our model described in Section 3, our proposed learning algorithm optimizes two rank loss functions alternately for joint visual and semantic embedding learning. With the regularized RankNet loss functions, we would exhibit the learning behavior during the training. As illustrated in Figure 4, the regularized rank losses, L_v and L_s , with respect to the visual and the semantic models keep decreasing steadily on training data as the training epochs increase regardless of the data split settings and datasets. Nevertheless, we adopt the early-stop strategy to avoid overfitting. However, we observe that the change of two ranking losses on validation data fluctuates wildly in learning so that we cannot decide a proper early-stop point easily. Instead we use the I-MAP measured on validation data to decide the proper early-stop points, as shown in Figure 4 where the bars of dash line indicate the actual point that the learning is stopped for different training datasets. In general, all our experiments in learning (including not shown in Figure 4) suggest that our alternate learning algorithm always converges regardless of different rank losses and datasets under different data split settings.

5.2. Results on Comparison to Baseline Models

Tables 4 and 5 summarize all the results yielded by four baseline models described in Section 4.6.1 and the full model described in Section 3.3, with the use of regularized RankNet loss and hinge rank loss described in Section 3.3.2 respectively. The experimental results are reported based on two different data split settings described in Section 4.1.2, *instance-first split* (IFS) and *label-first split* (LFS), under three different evaluation scenarios described in Section 4.4; i.e., *generalized ZSL*, *known-action only* and *unseen-action only* scenarios. For reliability, we report the *mean* and *standard error of the mean* (SEM) of results ($k = 5$ used in evaluation metrics, i.e., Eqs.(19-21)) over three randomly generated known/unseen label splits for each evaluation scenario.

For the IFS setting, it is observed from Tables 4 and 5 that all the baseline models and the full model perform significantly better than the RGS, a random guess model, on two datasets regardless of evaluation scenarios apart from the RLR model under the unseen-action only scenario. Due to a lack of knowledge transfer in a random label representation, the zero-shot performance of the RLR is expected. Overall, the full model outperforms all the baseline models on both datasets in the generalized ZSL and unseen-action only scenarios regardless of evaluation metrics. A comparison to the WSE suggests that the performance of the full model is generally superior to this baseline on both datasets under different evaluation scenarios, which lends evidence to support the necessity of the semantic embedding learning in multi-label learning problems. Also, we observe that the full model outperforms the RLR on Breakfast but fails

Table 3: Optimal hyper-parameter values of different learning models found by grid search. **Notation:** IFS – Instance-First Split; LFS – Label-First Split; V – Visual model; S – Semantic model; lr – learning rate; C, ϵ – soft-margin and percentage of support vectors in SVM/SVR; m – margin in the hinge ranking loss. $N_1 \rightarrow N_2 \rightarrow d_e$ indicates a neural network architecture where N_1 (dropout rate) is the number of neurons in the first hidden layer and dropout rate used in learning; N_2 is the number of hidden neurons in the second hidden layer; and d_e is the number of neurons in the latent embedding layer.

Dataset	Data Split	Model	Split		
			1	2	3
Breakfast	IFS	NRC(RankNet)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 2048(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500
		NRC(Hinge)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$
		WSE(RankNet)	V: lr = $1e-4$; 512(0) \rightarrow 2048	V: lr = $1e-4$; 512(0) \rightarrow 1024	V: lr = $1e-4$; 1024(0) \rightarrow 2048
		WSE(Hinge)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048; $m = 1$	V: lr = $1e-4$; 512(0) \rightarrow 1024; $m = 1$	V: lr = $1e-4$; 1024(0) \rightarrow 1024; $m = 1$
		RLR(RankNet)	V: lr = $1e-4$; 512(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 200 S: lr = $1e-6$; 300 \rightarrow 200
		RLR(Hinge)	V: lr = $1e-4$; 1024(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 200 S: lr = $1e-6$; 500 \rightarrow 200; $m = 10$
		DSP	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$
		ConSE	$C = 1$	$C = 1$	$C = 1$
		COSTA	$C = 1$	$C = 1$	$C = 1$
		Fast0Tag	V: lr = $1e-4$; 8192(0) \rightarrow 1024	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048	V: lr = $1e-2$; 8192(0.5) \rightarrow 2048
		Fast0Tag+	V: lr = $1e-4$; 4096(0) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 8192(0) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 300 \rightarrow 200
		Ours(RankNet)	V: lr = $1e-4$; 512(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 1024(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500	V: lr = $1e-4$; 1024(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500
		Ours(Hinge)	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 1024(0) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500; $m = 1$
	LFS	NRC(RankNet)	V: lr = $1e-4$; 2048(0.5) \rightarrow 2048 \rightarrow 200 S: lr = $1e-6$; 500 \rightarrow 200	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500
		NRC(Hinge)	V: lr = $1e-4$; 4096(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 300 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 4096(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 300 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 2048(0.5) \rightarrow 2048 \rightarrow 200 S: lr = $1e-6$; 300 \rightarrow 200; $m = 1$
		WSE(RankNet)	V: lr = $1e-4$; 1024(0) \rightarrow 1024	V: lr = $1e-4$; 1024(0) \rightarrow 1024	V: lr = $1e-4$; 1024(0.5) \rightarrow 1024
		WSE(Hinge)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048; $m = 10$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048; $m = 10$	V: lr = $1e-4$; 512(0) \rightarrow 2048; $m = 1$
		RLR(RankNet)	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 700 \rightarrow 200	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 200 S: lr = $1e-6$; 500 \rightarrow 200
		RLR(Hinge)	V: lr = $1e-4$; 512(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 256(0) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 512(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 10$
		DSP	$C = 100, \epsilon = 0.1$	$C = 100, \epsilon = 0.1$	$C = 100, \epsilon = 0.1$
		ConSE	$C = 100$	$C = 100$	$C = 100$
		COSTA	$C = 100$	$C = 100$	$C = 100$
		Fast0Tag	V: lr = $1e-4$; 4096(0.5) \rightarrow 2048	V: lr = $1e-4$; 4096(0) \rightarrow 1024	V: lr = $1e-2$; 8192(0.5) \rightarrow 2048
		Fast0Tag+	V: lr = $1e-4$; 8192(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 8192(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 4096(0) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 300 \rightarrow 200
		Ours(RankNet)	V: lr = $1e-4$; 512(0) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 256(0) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 700 \rightarrow 200
		Ours(Hinge)	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 500 \rightarrow 200; $m = 10$	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 256(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500; $m = 1$
Charades	IFS	NRC(RankNet)	V: lr = $1e-4$; 1024(0) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 1024(0) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 2048(0) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 700 \rightarrow 200
		NRC(Hinge)	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 4096(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 300 \rightarrow 500; $m = 10$
		WSE(RankNet)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048	V: lr = $1e-4$; 512(0.5) \rightarrow 2048	V: lr = $1e-4$; 256(0.5) \rightarrow 2048
		WSE(Hinge)	V: lr = $1e-4$; 1024(0) \rightarrow 2048	V: lr = $1e-4$; 1024(0) \rightarrow 1024	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048
		RLR(RankNet)	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500
		RLR(Hinge)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 10$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 10$
		DSP	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$
		ConSE	$C = 1$	$C = 1$	$C = 1$
		COSTA	$C = 1$	$C = 1$	$C = 1$
		Fast0Tag	V: lr = $1e-4$; 4096(0.5) \rightarrow 2048	V: lr = $1e-4$; 8192(0) \rightarrow 2048	V: lr = $1e-4$; 4096(0.5) \rightarrow 2048
		Fast0Tag+	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 8192(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800
		Ours(RankNet)	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 1024(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800
		Ours(Hinge)	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 1$	V: lr = $1e-4$; 1024(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 1$
	LFS	NRC(RankNet)	V: lr = $1e-4$; 1024(0) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 2048(0) \rightarrow 1024 \rightarrow 200 S: lr = $1e-6$; 700 \rightarrow 200	V: lr = $1e-4$; 4096(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 300 \rightarrow 500
		NRC(Hinge)	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 2048(0) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 200; $m = 1$	V: lr = $1e-4$; 2048(0) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$
		WSE(RankNet)	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048	V: lr = $1e-4$; 512(0.5) \rightarrow 2048	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048
		WSE(Hinge)	V: lr = $1e-4$; 1024(0) \rightarrow 2048; $m = 10$	V: lr = $1e-4$; 256(0.5) \rightarrow 2048; $m = 1$	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048; $m = 10$
		RLR(RankNet)	V: lr = $1e-4$; 256(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800
		RLR(Hinge)	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800; $m = 10$	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 10$
		DSP	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$	$C = 1, \epsilon = 0.1$
		ConSE	$C = 1$	$C = 1$	$C = 1$
		COSTA	$C = 1$	$C = 1$	$C = 1$
		Fast0Tag	V: lr = $1e-4$; 4096(0.5) \rightarrow 1024	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048	V: lr = $1e-4$; 4096(0.5) \rightarrow 1024
		Fast0Tag+	V: lr = $1e-4$; 8192(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 8192(0) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500	V: lr = $1e-4$; 4096(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 500 \rightarrow 500
		Ours(RankNet)	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 700 \rightarrow 800	V: lr = $1e-4$; 512(0.5) \rightarrow 2048 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800	V: lr = $1e-4$; 1024(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500
		Ours(Hinge)	V: lr = $1e-4$; 512(0.5) \rightarrow 1024 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 256(0.5) \rightarrow 2048 \rightarrow 500 S: lr = $1e-6$; 700 \rightarrow 500; $m = 1$	V: lr = $1e-4$; 1024(0.5) \rightarrow 1024 \rightarrow 800 S: lr = $1e-6$; 500 \rightarrow 800; $m = 1$

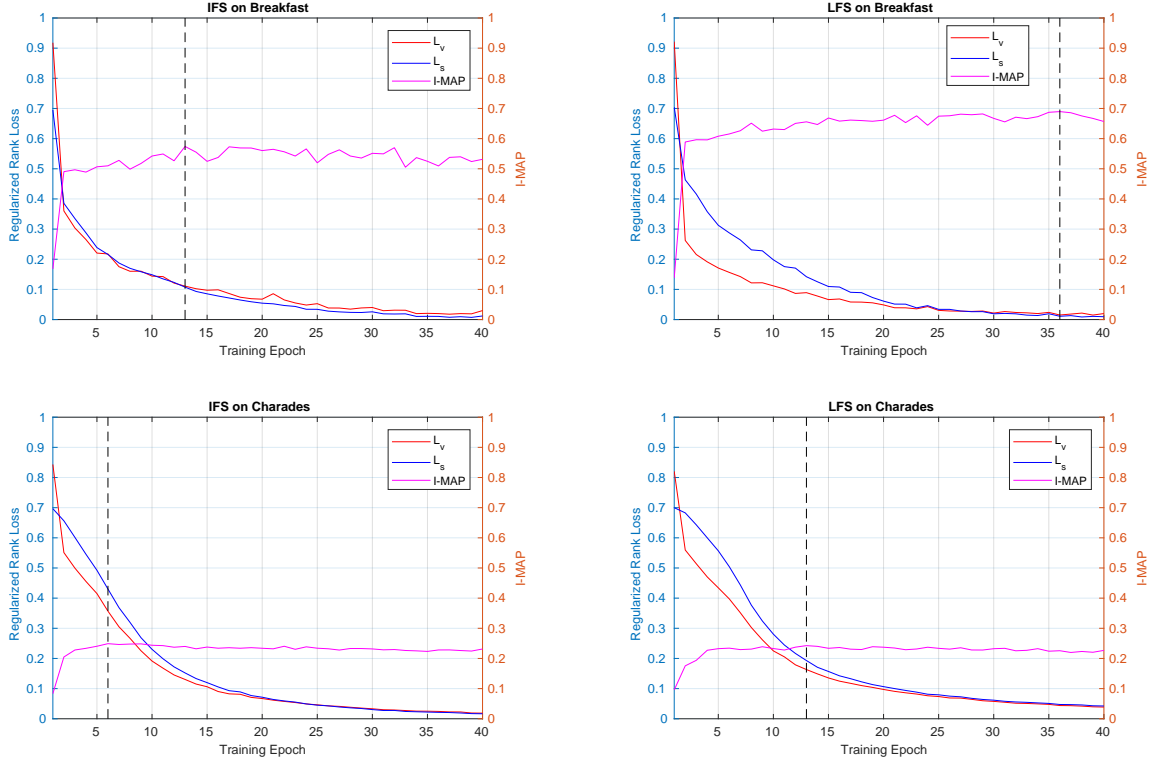


Figure 4: The evolution of the regularized RankNet losses, L_v and L_s , on training data and the I-MAP values on validation data during the joint visual and semantic embedding learning with our alternate learning algorithm described in Algorithm 1. All the results are achieved based on Split 1 on two datasets under the IFS and the LFS settings (c.f. Table 2).

Table 4: Multi-label zero-shot recognition performance (mean \pm SEM%) of the baseline models and the full model using **RankNet loss** in three evaluation scenarios under different data split settings. **Notation:** GZSL – generalized ZSL scenario; KnownA – Known-action only scenario; UnseenA – Unseen-action only scenario. Other notations are the same as described in Table 3.

Data Split	Evaluation Scenario	Model	Breakfast					Charades				
			L-MAP	I-MAP	P	R	F ₁	L-MAP	I-MAP	P	R	F ₁
IFS	GZSL	RGS	10.9 \pm 0.0	15.4 \pm 0.0	8.8 \pm 0.0	10.2 \pm 0.0	9.4 \pm 0.0	5.9 \pm 0.0	8.5 \pm 0.0	5.6 \pm 0.0	3.2 \pm 0.0	4.1 \pm 0.0
		NRC	27.9 \pm 0.7	49.1 \pm 1.1	35.1 \pm 0.9	40.5 \pm 1.0	37.6 \pm 1.0	9.2 \pm 0.1	20.9 \pm 0.7	20.7 \pm 0.9	11.9 \pm 0.5	15.1 \pm 0.7
		WSE	30.2 \pm 0.4	50.1 \pm 0.4	36.7 \pm 0.6	42.4 \pm 0.7	39.3 \pm 0.6	9.3 \pm 0.1	20.6 \pm 0.4	20.5 \pm 0.5	11.7 \pm 0.3	14.9 \pm 0.3
		RRL	29.6 \pm 0.3	50.4 \pm 0.5	36.7 \pm 0.7	42.3 \pm 0.8	39.3 \pm 0.8	9.4 \pm 0.1	20.7 \pm 1.0	21.1 \pm 1.4	12.0 \pm 0.8	15.3 \pm 1.0
		Ours	32.8\pm0.7	53.5\pm1.2	38.6\pm1.6	44.5\pm1.9	41.4\pm1.8	9.7\pm0.1	22.4\pm0.4	22.7\pm0.4	13.0\pm0.2	16.5\pm0.3
	KnownA	RGS	11.4 \pm 0.2	17.4 \pm 0.1	9.6 \pm 0.2	12.9 \pm 0.0	11.0 \pm 0.1	6.1 \pm 0.1	9.4 \pm 0.1	5.8 \pm 0.1	4.3 \pm 0.0	4.9 \pm 0.0
		NRC	30.3 \pm 1.0	53.6 \pm 0.8	35.4 \pm 0.7	47.5 \pm 0.6	40.6 \pm 0.6	10.0 \pm 0.2	25.4 \pm 0.5	23.0 \pm 0.6	17.1 \pm 0.2	19.6 \pm 0.4
		WSE	32.1 \pm 0.9	55.4 \pm 1.0	37.6 \pm 0.3	50.6 \pm 0.5	43.2 \pm 0.2	10.1 \pm 0.3	24.6 \pm 0.3	22.3 \pm 0.3	16.6 \pm 0.0	19.0 \pm 0.1
		RRL	33.5 \pm 1.1	56.3 \pm 0.8	37.4 \pm 1.0	50.2 \pm 0.8	42.8 \pm 0.9	10.7\pm0.1	26.7\pm0.7	23.9\pm0.9	17.8\pm0.4	20.4\pm0.6
		Ours	35.3\pm1.3	58.0\pm0.4	38.2\pm1.6	51.3\pm1.4	43.8\pm1.5	10.5 \pm 0.2	26.1 \pm 0.3	23.5 \pm 0.5	17.5 \pm 0.1	20.0 \pm 0.2
	UnseenA	RGS	8.5 \pm 0.6	30.7 \pm 0.6	6.1 \pm 0.6	50.0 \pm 0.0	10.9 \pm 0.9	5.4 \pm 0.0	13.9 \pm 0.0	5.1 \pm 0.0	12.5 \pm 0.0	7.2 \pm 0.0
		NRC	17.1 \pm 0.7	47.5 \pm 2.4	8.7 \pm 0.8	70.8 \pm 2.5	15.4 \pm 1.2	6.8 \pm 0.4	20.4 \pm 1.1	8.4 \pm 0.7	20.8 \pm 1.9	11.9 \pm 1.0
		WSE	21.7 \pm 1.7	47.7 \pm 3.0	9.1 \pm 1.9	72.4 \pm 8.8	16.1 \pm 3.2	6.9 \pm 0.4	20.5 \pm 1.6	8.5 \pm 0.7	21.2 \pm 2.1	12.1 \pm 1.0
		RRL	13.2 \pm 2.5	34.3 \pm 7.1	7.0 \pm 1.6	56.6 \pm 9.7	12.4 \pm 2.7	5.5 \pm 0.4	13.4 \pm 0.6	4.7 \pm 0.4	11.6 \pm 0.6	6.7 \pm 0.5
		Ours	22.3\pm1.1	53.1\pm5.8	9.5\pm1.5	77.2\pm8.4	16.9\pm2.5	7.1\pm0.4	22.4\pm2.1	9.5\pm1.0	23.5\pm2.7	13.5\pm1.5
LFS	GZSL	RGS	15.2 \pm 1.5	17.8 \pm 0.4	11.3 \pm 0.4	10.2 \pm 0.0	10.7 \pm 0.2	5.2 \pm 0.1	7.9 \pm 0.1	5.1 \pm 0.1	3.2 \pm 0.0	3.9 \pm 0.0
		NRC	21.9 \pm 2.2	27.6 \pm 1.2	19.2 \pm 1.0	17.3 \pm 1.1	18.2 \pm 1.0	8.9 \pm 0.2	20.5 \pm 0.5	21.2 \pm 0.8	13.3 \pm 0.5	16.4 \pm 0.6
		WSE	25.2\pm1.4	31.0 \pm 2.7	23.0 \pm 2.0	20.7 \pm 1.4	21.7 \pm 1.6	9.1 \pm 0.0	20.3 \pm 0.0	21.0 \pm 0.2	13.2 \pm 0.2	16.2 \pm 0.2
		RRL	22.6 \pm 2.1	29.6 \pm 1.9	22.0 \pm 0.7	19.8 \pm 0.2	20.9 \pm 0.3	9.3\pm0.2	19.6 \pm 0.4	19.8 \pm 0.6	12.4 \pm 0.4	15.2 \pm 0.5
		Ours	25.0 \pm 1.4	32.6\pm2.9	23.6\pm2.1	21.2\pm1.5	22.3\pm1.7	9.2 \pm 0.1	20.8\pm0.3	21.6\pm0.7	13.5\pm0.4	16.6\pm0.5
	KnownA	RGS	14.7 \pm 1.9	17.8 \pm 0.2	10.1 \pm 0.2	12.8 \pm 0.0	11.3 \pm 0.1	5.0 \pm 0.1	8.0 \pm 0.1	4.8 \pm 0.1	3.7 \pm 0.0	4.2 \pm 0.0
		NRC	22.2 \pm 2.5	30.2 \pm 1.5	18.9 \pm 1.1	24.2 \pm 1.8	21.2 \pm 1.4	8.9 \pm 0.3	22.3 \pm 0.6	21.1 \pm 0.7	15.9 \pm 0.4	18.1 \pm 0.5
		WSE	25.3\pm2.0	34.2 \pm 3.3	23.1 \pm 1.9	29.5 \pm 2.7	25.9 \pm 2.3	9.0 \pm 0.0	22.2 \pm 0.3	21.1 \pm 0.4	15.9 \pm 0.0	18.1 \pm 0.1
		RRL	24.8 \pm 2.6	34.3 \pm 3.1	21.7 \pm 0.7	27.6 \pm 0.4	24.3 \pm 0.6	9.5\pm0.1	23.0\pm0.4	21.2 \pm 0.5	16.0 \pm 0.4	18.2 \pm 0.4
		Ours	25.2 \pm 1.6	35.7\pm3.6	23.3\pm2.0	29.6\pm2.8	26.0\pm2.3	9.1 \pm 0.1	23.0\pm0.6	21.6\pm0.7	16.3\pm0.3	18.6\pm0.5
	UnseenA	RGS	16.8 \pm 1.6	34.2 \pm 1.3	16.2 \pm 1.6	50.0 \pm 0.0	24.4 \pm 1.8	6.9 \pm 0.1	19.4 \pm 0.1	6.7 \pm 0.1	25.0 \pm 0.0	10.5 \pm 0.1
		NRC	21.2 \pm 1.4	42.3 \pm 5.4	18.7 \pm 3.2	57.1 \pm 7.0	28.0 \pm 4.4	9.4 \pm 0.1	30.7\pm2.2	11.6\pm0.7	43.2\pm2.8	18.2\pm1.1
		WSE	25.2\pm0.4	42.5 \pm 4.6	20.0 \pm 3.1	61.5 \pm 6.8	30.1 \pm 4.3	9.7\pm0.1	30.1 \pm 1.2	11.5 \pm 0.4	43.1 \pm 1.8	18.2\pm0.6
		RRL	16.7 \pm 2.5	32.2 \pm 2.1	15.0 \pm 1.9	45.8 \pm 2.0	22.5 \pm 2.4	7.7 \pm 0.5	19.3 \pm 0.3	6.7 \pm 0.1	25.2 \pm 0.6	10.6 \pm 0.2
		Ours	24.6 \pm 1.5	44.8\pm4.7	22.2\pm2.5	68.6\pm4.7	33.4\pm3.3	9.7\pm0.1	29.2 \pm 2.1	11.1 \pm 0.9	41.6 \pm 3.5	17.5 \pm 1.4

Table 5: Multi-label zero-shot recognition performance (mean \pm SEM%) of the baseline models and our full model using **Hinge rank loss** in three evaluation scenarios under different data split settings. **Notation:** GZSL – generalized ZSL scenario; KnownA – Known-action only scenario; UnseenA – Unseen-action only scenario. Other notations are the same as described in Table 3.

Data Split	Evaluation Scenario	Model	Breakfast					Charades				
			L-MAP	I-MAP	P	R	F ₁	L-MAP	I-MAP	P	R	F ₁
IFS	GZSL	RGS	10.9 \pm 0.0	15.4 \pm 0.0	8.8 \pm 0.0	10.2 \pm 0.0	9.4 \pm 0.0	5.9 \pm 0.0	8.5 \pm 0.0	5.6 \pm 0.0	3.2 \pm 0.0	4.1 \pm 0.0
		NRC	28.1 \pm 0.4	50.1 \pm 1.0	36.7 \pm 0.7	42.3 \pm 0.8	39.3 \pm 0.7	9.2 \pm 0.1	21.7 \pm 0.7	22.2 \pm 0.7	12.7 \pm 0.4	16.2 \pm 0.5
		WSE	31.1 \pm 0.7	50.4 \pm 0.3	36.6 \pm 0.5	42.2 \pm 0.6	39.2 \pm 0.6	8.6 \pm 0.0	20.1 \pm 0.5	20.3 \pm 0.7	11.6 \pm 0.4	14.7 \pm 0.5
		RLR	30.3 \pm 1.2	51.8 \pm 1.8	37.4 \pm 1.4	43.2 \pm 1.6	40.1 \pm 1.5	9.5 \pm 0.1	21.4 \pm 0.9	22.4 \pm 0.8	12.8 \pm 0.5	16.3 \pm 0.6
		Ours	32.7\pm0.4	53.4\pm0.8	38.9\pm0.5	44.9\pm0.5	41.7\pm0.5	10.0\pm0.1	22.6\pm0.4	23.1\pm0.5	13.2\pm0.3	16.8\pm0.4
	KnownA	RGS	11.4 \pm 0.2	17.4 \pm 0.1	9.6 \pm 0.2	12.9 \pm 0.0	11.0 \pm 0.1	6.1 \pm 0.1	9.4 \pm 0.1	5.8 \pm 0.1	4.3 \pm 0.0	4.9 \pm 0.0
		NRC	30.5 \pm 0.6	54.9 \pm 0.9	36.6 \pm 0.4	49.3 \pm 0.4	42.0 \pm 0.2	10.1 \pm 0.3	25.4 \pm 0.5	23.2 \pm 0.6	17.2 \pm 0.2	19.8 \pm 0.4
		WSE	33.0 \pm 0.9	55.4 \pm 0.7	37.1 \pm 0.6	49.9 \pm 0.6	42.6 \pm 0.5	9.3 \pm 0.2	23.7 \pm 0.4	21.5 \pm 0.5	16.0 \pm 0.2	18.3 \pm 0.3
		RLR	34.9 \pm 1.8	59.0\pm1.3	38.6\pm0.7	51.9 \pm 1.1	44.3\pm0.8	10.8 \pm 0.2	26.5\pm0.7	23.7 \pm 0.8	17.7 \pm 0.4	20.3 \pm 0.5
		Ours	35.6\pm0.5	58.2 \pm 1.0	38.6\pm0.3	52.0\pm0.5	44.3\pm0.0	10.9\pm0.2	26.4 \pm 0.3	24.1\pm0.3	17.9\pm0.2	20.6\pm0.1
	UnseenA	RGS	8.5 \pm 0.6	30.7 \pm 0.6	6.1 \pm 0.6	50.0 \pm 0.0	10.9 \pm 0.9	5.4 \pm 0.0	13.9 \pm 0.0	5.1 \pm 0.0	12.5 \pm 0.0	7.2 \pm 0.0
		NRC	17.6 \pm 1.6	46.7 \pm 3.2	9.1\pm1.2	73.4\pm4.0	16.1\pm2.0	6.7 \pm 0.4	21.8 \pm 1.2	9.3 \pm 0.4	23.1 \pm 1.4	13.2 \pm 0.6
		WSE	22.7\pm2.1	43.6 \pm 6.0	9.1\pm1.9	72.6 \pm 8.3	16.1\pm3.1	6.6 \pm 0.4	21.3 \pm 1.9	8.7 \pm 0.7	21.7 \pm 2.6	12.4 \pm 1.1
		RLR	10.5 \pm 2.0	35.5 \pm 2.1	5.9 \pm 0.3	48.4 \pm 2.1	10.5 \pm 0.5	5.6 \pm 0.3	14.7 \pm 1.9	5.0 \pm 0.7	12.4 \pm 2.0	7.1 \pm 1.1
		Ours	20.3 \pm 0.6	47.6\pm1.6	8.8 \pm 1.0	71.6 \pm 2.2	15.7 \pm 1.6	7.3\pm0.5	22.6\pm1.5	9.6\pm0.7	23.9\pm2.0	13.7\pm1.0
LFS	GZSL	RGS	15.2 \pm 1.5	17.8 \pm 0.4	11.3 \pm 0.4	10.2 \pm 0.0	10.7 \pm 0.2	5.2 \pm 0.1	7.9 \pm 0.1	5.1 \pm 0.1	3.2 \pm 0.0	3.9 \pm 0.0
		NRC	22.6 \pm 2.0	26.6 \pm 1.7	18.5 \pm 1.3	16.7 \pm 1.2	17.5 \pm 1.2	8.9 \pm 0.1	20.7 \pm 0.1	21.2 \pm 0.3	13.3 \pm 0.3	16.4 \pm 0.3
		WSE	24.4 \pm 1.7	31.3 \pm 3.0	23.4 \pm 1.9	21.1 \pm 1.4	22.2 \pm 1.6	8.1 \pm 0.2	19.4 \pm 0.2	19.4 \pm 0.4	12.2 \pm 0.3	14.9 \pm 0.3
		RLR	22.2 \pm 1.4	31.6 \pm 1.6	24.3 \pm 2.1	21.8 \pm 1.6	23.0 \pm 1.8	9.0 \pm 0.1	19.9 \pm 0.7	20.7 \pm 1.0	13.0 \pm 0.5	16.0 \pm 0.6
		Ours	24.7\pm1.4	32.9\pm2.9	24.6\pm2.5	22.1\pm1.9	23.3\pm2.2	9.2\pm0.1	21.1\pm0.4	22.1\pm0.7	13.9\pm0.4	17.1\pm0.5
	KnownA	RGS	14.7 \pm 1.9	17.8 \pm 0.2	10.1 \pm 0.2	12.8 \pm 0.0	11.3 \pm 0.1	5.0 \pm 0.1	8.0 \pm 0.1	4.8 \pm 0.1	3.7 \pm 0.0	4.2 \pm 0.0
		NRC	23.0 \pm 2.6	29.4 \pm 2.7	18.5 \pm 1.5	23.6 \pm 2.2	20.7 \pm 1.8	8.8 \pm 0.1	22.6 \pm 0.4	21.3 \pm 0.4	16.1 \pm 0.0	18.3 \pm 0.1
		WSE	24.9 \pm 2.4	35.1 \pm 3.8	23.4 \pm 2.0	29.7 \pm 2.0	26.2 \pm 2.0	7.9 \pm 0.2	21.1 \pm 0.5	19.6 \pm 0.4	14.8 \pm 0.1	16.9 \pm 0.2
		RLR	23.6 \pm 2.6	36.1 \pm 2.7	24.3 \pm 2.2	31.0 \pm 3.2	27.2 \pm 2.6	9.2\pm0.1	22.7 \pm 0.9	21.0 \pm 1.0	15.9 \pm 0.5	18.1 \pm 0.7
		Ours	25.2\pm1.8	37.1\pm3.4	25.0\pm2.0	31.9\pm2.7	28.1\pm2.3	9.1 \pm 0.0	23.2\pm0.4	22.1\pm0.5	16.6\pm0.2	19.0\pm0.3
	UnseenA	RGS	16.8 \pm 1.6	34.2 \pm 1.3	16.2 \pm 1.6	50.0 \pm 0.0	24.4 \pm 1.8	6.9 \pm 0.1	19.4 \pm 0.1	6.7 \pm 0.1	25.0 \pm 0.0	10.5 \pm 0.1
		NRC	21.6 \pm 0.6	39.6 \pm 2.8	19.7\pm1.4	61.4\pm4.6	29.7\pm1.8	9.5 \pm 0.2	30.7 \pm 2.3	11.8 \pm 0.7	44.0 \pm 2.8	18.6 \pm 1.1
		WSE	23.3\pm0.2	41.2\pm4.3	19.0 \pm 2.8	58.4 \pm 5.7	28.6 \pm 3.8	9.3 \pm 0.1	29.9 \pm 1.6	11.5 \pm 0.2	42.9 \pm 1.3	18.1 \pm 0.4
		RLR	18.9 \pm 1.1	33.3 \pm 1.2	16.4 \pm 0.1	51.5 \pm 5.4	24.7 \pm 0.6	7.4 \pm 0.5	18.5 \pm 0.8	6.4 \pm 0.5	23.8 \pm 1.4	10.1 \pm 0.7
		Ours	23.3\pm2.0	40.2 \pm 5.1	19.4 \pm 3.6	59.0 \pm 7.7	29.1 \pm 5.0	10.0\pm0.1	30.9\pm2.6	12.0\pm0.9	44.9\pm3.8	18.9\pm1.5

to do so on Charades in the know-action only scenario when using RankNet loss (Table 4). This result suggests that a semantic representation of labels is not critically important for known actions in multi-label learning when the semantic embedding learning has been employed to explore the between-action relations from label co-occurrences. This observation further implies that the semantic embedding learning cannot explore the semantic relations between labels properly unless there are sufficient training examples for different actions. Nevertheless, the performance in the unseen-action only and the generalized ZSL scenarios clearly indicates the importance of the semantic representation of an action label for knowledge transfer required by ZSL. It is also evident from Tables 4 and 5 that the full model always outperforms the NRC where there are no recurrent connections. Thus, the comparison to the baseline models clearly suggest that the performance gain is brought by the use of an LSTM layer in the visual model and the semantic embedding learning fulfilled in the semantic model.

For the LFS setting, results shown in Tables 4 and 5 suggest that all the baseline models perform significantly better than a random guess. Overall, the full model generally outperforms those baseline models on both datasets. In few circumstances, however, the full model slightly under-performs the WSE on Breakfast in terms of L-MAP with a tiny margin (Table 4). Besides, it is observed from Table 4 that in the unseen-action only scenario, our model slightly under-performs the WSE and the NRC on Charades although it yields the best performance on Breakfast. While from Table 5 we can observe that our full model performs the best on Charades but not on Breakfast. These results reveal that the two employed rank losses are complementary when learning the joint embedding space.

In summary, the comparison to the elaborated baseline models facilitates the understanding of different components and ranking loss functions employed in our proposed framework for multi-label zero-shot human action recognition. Two different ranking losses used in our framework yield the similar performance overall. By comparison to four baseline models, the full model generally leads to better results on two datasets measured with different evaluation metrics in all three evaluation scenarios, although the experimental results also reveal the limitation of components used in the full model to be investigated in our future studies.

5.3. Results on Comparison to State-of-the-Art Methods

Table 6 summarizes the experimental results of the comparative study described in Section 4.6.2. Multi-label ZSL performance of five different methods including Fast0Tag+ (our extension for Fast0Tag) with the reference to a random guess is reported to be compared with our proposed framework where two different rank losses and their fusion are employed, respectively. Again, all the experiments are conducted with two different data split settings and evaluated under three evaluation scenarios, as described in Section 5.2. For reliability, we report the mean and the SEM of results ($k = 5$ used in evaluation metrics, i.e., Eqs(19-21)) over three randomly generated known/unseen label splits under each evaluation scenario.

For the IFS setting, it is seen from Table 6 that all the models perform better than random guess in most of evaluation scenarios. However, DSP and ConSE result in the poorer performance than random guess in the unseen-action only scenario on Breakfast in terms of some specific metric, e.g., I-MAP. Overall, DSP and ConSE under-perform other methods considerably in terms

Table 6: Multi-label zero-shot recognition performance (mean \pm SEM%) of five state-of-the-art methods and ours with the reference to random guess in three evaluation scenarios under different data split settings. The notations are the same as used in Tables 3-5.

Data Split	Evaluation Scenario	Model	Breakfast					Charades				
			L-MAP	I-MAP	P	R	F ₁	L-MAP	I-MAP	P	R	F ₁
IFS	GZSL	RGS	10.9±0.0	15.4±0.0	8.8±0.0	10.2±0.0	9.4±0.0	5.9±0.0	8.5±0.0	5.6±0.0	3.2±0.0	4.1±0.0
		DSP[16]	21.3±0.0	25.4±0.9	16.8±0.4	19.4±0.4	18.0±0.4	7.9±0.0	12.5±0.1	12.5±0.2	7.2±0.1	9.1±0.1
		ConSE[62]	16.1±0.2	29.6±0.3	20.3±0.1	23.4±0.1	21.7±0.1	7.3±0.0	13.9±0.3	14.6±0.5	8.3±0.3	10.6±0.4
		COSTA[22]	19.7±0.2	37.4±0.3	28.8±0.4	33.2±0.5	30.8±0.5	8.5±0.1	16.9±0.7	17.6±0.9	10.1±0.5	12.8±0.6
		FastOTag[6]	22.3±1.1	38.6±0.2	27.5±0.1	31.8±0.1	29.5±0.1	9.3±0.0	20.6±1.0	20.6±1.4	11.8±0.8	15.0±1.0
		FastOTag+	23.4±0.3	40.6±0.7	29.2±0.3	33.7±0.3	31.3±0.3	9.7±0.1	21.7±0.6	21.9±0.8	12.5±0.5	15.9±0.6
		Ours(RankNet)	32.1±0.8	53.3±1.0	39.0±0.9	45.0±1.0	41.8±0.9	9.7±0.1	22.4±0.4	22.8±0.4	13.0±0.2	16.5±0.3
		Ours(Hinge)	32.7±0.4	53.4±0.8	38.9±0.5	44.9±0.5	41.7±0.5	10.0±0.1	22.6±0.4	23.1±0.5	13.2±0.3	16.8±0.4
		Ours(Fusion)	33.9±0.4	54.7±1.1	40.0±1.1	46.1±1.3	42.8±1.2	10.1±0.1	23.3±0.5	23.7±0.7	13.5±0.4	17.2±0.5
	KnownA	RGS	11.4±0.2	17.4±0.1	9.6±0.2	12.9±0.0	11.0±0.1	6.1±0.1	9.4±0.1	5.8±0.1	4.3±0.0	4.9±0.0
		DSP[16]	22.6±0.5	30.8±1.9	19.1±0.9	25.7±1.3	21.9±1.1	8.5±0.1	13.9±0.3	12.4±0.4	9.2±0.3	10.6±0.3
		ConSE[62]	17.1±0.3	33.0±1.2	21.0±0.2	28.2±0.7	24.1±0.4	7.8±0.2	15.9±0.4	14.6±0.7	10.8±0.3	12.4±0.5
		COSTA[22]	22.1±0.7	41.6±0.7	28.8±0.4	38.7±0.2	33.0±0.3	9.3±0.2	19.7±0.6	17.6±0.9	13.1±0.5	15.0±0.6
		FastOTag[6]	23.9±1.3	44.3±0.5	29.8±0.7	40.0±0.6	34.1±0.6	10.0±0.2	24.6±0.6	22.4±0.7	16.6±0.3	19.1±0.4
		FastOTag+	25.4±0.2	45.0±0.5	29.9±0.4	40.2±1.1	34.3±0.6	10.5±0.3	25.7±0.7	23.3±0.9	17.3±0.4	19.8±0.6
		Ours(RankNet)	34.5±1.0	57.8±0.7	38.5±1.0	51.8±0.6	44.2±0.9	10.5±0.2	26.1±0.3	23.5±0.5	17.5±0.1	20.0±0.2
		Ours(Hinge)	35.6±0.5	58.2±1.0	38.6±0.3	52.0±0.5	44.3±0.0	10.9±0.2	26.4±0.3	24.1±0.3	17.9±0.2	20.6±0.1
		Ours(Fusion)	36.6±0.7	59.4±0.5	39.6±0.9	53.2±0.5	45.4±0.8	11.0±0.3	27.1±0.4	24.6±0.5	18.3±0.2	21.0±0.3
	UnseenA	RGS	8.5±0.6	30.7±0.6	6.1±0.6	50.0±0.0	10.9±0.9	5.4±0.0	13.9±0.0	5.1±0.0	12.5±0.0	7.2±0.0
		DSP[16]	15.9±1.5	27.0±4.7	5.8±1.6	47.6±11.5	10.4±2.7	6.2±0.4	17.3±1.4	7.4±0.8	18.4±1.9	10.6±1.1
		ConSE[62]	12.2±0.3	29.9±5.0	6.2±0.9	51.0±6.9	11.1±1.5	5.8±0.4	17.3±0.8	7.0±0.7	17.2±1.1	9.9±0.8
		COSTA[22]	9.2±1.2	37.4±2.8	7.4±0.7	60.1±2.5	13.1±1.0	6.0±0.3	15.5±1.0	6.3±0.8	15.4±1.5	8.9±1.0
		FastOTag[6]	15.3±0.9	36.7±4.1	7.0±1.4	55.9±6.7	12.4±2.3	7.1±0.4	20.2±2.4	8.3±0.8	20.8±2.7	11.9±1.2
		FastOTag+	15.1±1.2	39.4±1.3	7.4±0.9	60.1±3.6	13.2±1.5	7.3±0.4	19.3±0.5	8.1±0.3	20.0±0.6	11.5±0.3
		Ours(RankNet)	21.9±0.3	51.0±4.5	9.4±1.2	76.5±6.8	16.7±2.0	7.1±0.4	22.4±2.1	9.5±1.0	23.5±2.7	13.5±1.5
		Ours(Hinge)	20.3±0.6	47.6±1.6	8.8±1.0	71.6±2.2	15.7±1.6	7.3±0.5	22.6±1.5	9.6±0.7	23.9±2.0	13.7±1.0
		Ours(Fusion)	22.3±0.4	52.9±4.6	9.7±1.6	78.8±6.8	17.3±2.6	7.3±0.5	23.1±1.8	9.9±0.8	24.6±2.5	14.1±1.2
LFS	GZSL	RGS	15.2±1.5	17.8±0.4	11.3±0.4	10.2±0.0	10.7±0.2	5.2±0.1	7.9±0.1	5.1±0.1	3.2±0.0	3.9±0.0
		DSP[16]	20.7±1.7	18.6±1.9	11.0±1.5	9.9±1.0	10.4±1.2	7.4±0.1	12.1±0.3	12.4±0.6	7.7±0.3	9.5±0.4
		ConSE[62]	18.5±2.1	20.2±1.8	12.7±1.1	11.4±0.6	12.0±0.8	7.0±0.1	13.8±0.1	14.8±0.1	9.3±0.1	11.4±0.1
		COSTA[22]	19.3±2.1	22.7±2.1	16.8±1.0	15.1±0.7	15.9±0.8	8.9±0.1	17.3±0.1	18.6±0.1	11.7±0.2	14.4±0.1
		FastOTag[6]	22.5±1.5	24.3±1.7	16.2±0.6	14.6±0.1	15.4±0.3	8.6±0.1	20.1±0.4	20.1±0.9	12.6±0.6	15.5±0.7
		FastOTag+	21.9±1.1	23.3±1.1	15.3±0.2	13.8±0.6	14.5±0.4	9.0±0.1	20.9±0.3	21.4±0.6	13.5±0.4	16.5±0.5
		Ours(RankNet)	25.0±1.4	32.6±2.9	23.6±2.1	21.2±1.5	22.3±1.7	9.2±0.1	20.8±0.3	21.6±0.7	13.5±0.4	16.6±0.5
		Ours(Hinge)	24.7±1.4	32.9±2.9	24.6±2.5	22.1±1.9	23.3±2.2	9.2±0.1	21.1±0.4	22.1±0.7	13.9±0.4	17.1±0.5
		Ours(Fusion)	25.5±1.4	33.3±2.5	24.6±2.3	22.1±1.6	23.2±1.9	9.6±0.1	21.5±0.4	22.6±0.6	14.2±0.4	17.4±0.4
	KnownA	RGS	14.7±1.9	17.8±0.2	10.1±0.2	12.8±0.0	11.3±0.1	5.0±0.1	8.0±0.1	4.8±0.1	3.7±0.0	4.2±0.0
		DSP[16]	14.1±1.2	18.1±2.7	8.5±1.8	10.3±2.1	9.4±1.9	7.3±0.1	12.9±0.4	12.3±0.7	9.3±0.3	10.6±0.5
		ConSE[62]	12.2±0.8	22.1±0.6	13.1±0.6	14.6±1.2	13.8±0.9	6.8±0.1	15.0±0.1	14.7±0.2	11.1±0.1	12.6±0.1
		COSTA[22]	19.8±2.9	25.1±2.7	16.8±1.0	21.3±0.8	18.7±0.9	9.0±0.1	19.5±0.1	18.6±0.1	14.1±0.2	16.0±0.1
		FastOTag[6]	23.1±2.0	26.5±1.9	15.9±0.4	20.2±0.8	17.8±0.5	8.4±0.1	22.4±0.7	20.8±0.7	15.7±0.2	17.9±0.4
		FastOTag+	22.2±1.7	24.9±1.5	15.1±0.6	19.3±1.1	17.0±0.8	9.0±0.1	22.8±0.5	21.3±0.5	16.1±0.1	18.3±0.3
		Ours(RankNet)	25.2±1.6	35.7±3.6	23.3±2.0	29.6±2.8	26.0±2.3	9.1±0.1	23.0±0.6	21.6±0.7	16.3±0.3	18.6±0.5
		Ours(Hinge)	25.2±1.8	37.1±3.4	25.0±2.0	31.9±2.7	28.1±2.3	9.1±0.0	23.2±0.4	22.1±0.5	16.6±0.2	19.0±0.3
		Ours(Fusion)	25.8±1.8	37.4±2.7	24.6±2.0	31.3±2.8	27.6±2.3	9.5±0.1	23.8±0.6	22.7±0.7	17.1±0.3	19.5±0.4
	UnseenA	RGS	16.8±1.6	34.2±1.3	16.2±1.6	50.0±0.0	24.4±1.8	6.9±0.1	19.4±0.1	6.7±0.1	25.0±0.0	10.5±0.1
		DSP[16]	16.6±0.8	28.8±1.4	14.7±0.6	42.8±1.3	21.9±0.7	8.5±0.2	22.8±0.9	8.1±0.4	30.3±1.3	12.8±0.6
		ConSE[62]	13.7±1.7	33.0±0.7	14.7±0.5	50.9±4.6	22.7±0.2	7.9±0.4	23.2±0.4	8.5±0.0	31.9±0.3	13.5±0.1
		COSTA[22]	18.2±1.1	35.8±1.2	16.9±1.7	52.2±0.0	25.4±2.0	7.6±0.2	24.9±1.0	9.2±0.4	34.4±1.8	14.5±0.7
		FastOTag[6]	21.1±1.1	39.1±4.0	18.5±3.6	56.0±6.5	27.8±4.9	9.5±0.1	28.4±2.5	11.0±0.5	41.1±2.2	17.4±0.8
		FastOTag+	21.2±0.8	45.5±3.4	19.9±1.8	61.7±2.4	30.0±2.1	9.3±0.1	30.2±2.9	11.6±0.7	43.2±3.1	18.2±1.2
		Ours(RankNet)	24.6±1.5	44.8±4.7	22.2±2.5	68.6±4.7	33.4±3.3	9.7±0.1	29.2±2.1	11.1±0.9	41.6±3.5	17.5±1.4
		Ours(Hinge)	23.3±2.0	40.2±5.1	19.4±3.6	59.0±7.7	29.1±5.0	10.0±0.1	30.9±2.6	12.0±0.9	44.9±3.8	18.9±1.5
		Ours(Fusion)	24.7±1.5	42.6±6.1	19.6±3.9	59.7±8.6	29.5±5.4	10.2±0.1	31.1±2.7	12.2±0.9	45.6±3.7	19.2±1.4

of all five evaluation metrics under all three evaluation scenarios. Such results demonstrate that simply combining semantic representations of co-occurred multiple labels into one collective representation leads to catastrophic loss of semantic information, which is mainly responsible for the poor performance of DSP and ConSE in multi-label recognition. Fast0Tag generally outperforms COSTA on two datasets in terms of most of evaluation metrics. While COSTA learns a classifier for each label separately without considering a relationship among co-occurred labels, the consideration of such a relationship in Fast0Tag accounts for the better performance. By incorporating the semantic embedding learning into Fast0Tag, Fast0Tag+, our extension of Fast0Tag, constantly improves the performance of its original version in most circumstances on two datasets. Once again, this result lends us evidence to justify the necessity of semantic embedding learning used in our framework for zero-shot multi-label ZSL. In contrast, our model trained with either RankNet loss or the hinge rank loss generally outperforms all five models significantly in terms of five evaluation metrics under different evaluation scenarios on two datasets, as highlighted with bold-font in Table 6. By comparing our model to Fast0Tag+, we see three main differences between them as follows: *visual representations*, *network architectures* for the visual model and *loss functions*. Regarding visual representations, our model uses the segment-based visual features for an instance while Fast0Tag+ employs an instance-level holistic visual representation. For network architectures, we employ an LSTM layer with recurrent connections to capture temporal coherence among segments of a video clip while Fast0Tag+ simply uses a feed-forward network. As described in Section 3.3.2, we use an alternative loss function to that in Fast0Tag. Thus, those differences together leverage our performance gain over Fast0Tag+, which yet again lends us evidence to support our proposed framework. Finally, it is observed from Table 6 that in the IFS setting, the RankNet and the hinge rank losses perform differently on two datasets; the hinge rank loss generally outperforms the RankNet loss on both datasets with the exceptions of unseen action only scenarios on Breakfast. Nevertheless, the fusion of two models trained with different losses leads to the best performance in most circumstances as highlighted with bold-font in Table 6. Such results reveal that two losses behave quite differently and the diversity can be exploited via fusion, which provides useful information to develop more effective rank loss functions.

For the LFS setting, experimental results suggest that most of the models in question have similar behavior to that in the IFS setting, as shown in Table 6. Once again, DSP and ConSE generally perform worse than other models and even under-perform random guess on Charades in the unseen-action only scenario. While COSTA yields better performance than DSP and ConSE overall, it generally under-performs Fast0Tag, Fast0Tag+ and ours in all three evaluation scenarios. In the LFS setting, our model trained with different rank losses generally outperforms others in most circumstances except for the unseen-action only scenario on Breakfast where Fast0Tag+ performs better than ours marginally in terms of I-MAP. Regarding two rank losses used in our experiments, the hinge rank loss marginally outper-

forms the RankNet loss in most circumstances on two datasets. Once again, the fusion of results brought by two rank losses further improves the performance in most circumstances, which provides the further evidence on the complementary aspect of two different rank losses. As described in Section 5.2, the LFS setting is more challenging than the IFS setting and some salient visual features on test instances corresponding to unseen actions could completely miss in training examples. In this case, the use of a segment-level based visual representation and an LSTM layer in the visual model may not be able to generalize well due to a lack of training examples. Although such a result does not sufficiently favour the use of a segment-level based visual representation and an LSTM layer in the visual model in the presence of limited training data, it is no doubt that introducing a semantic model to Fast0Tag leverages the performance gain. Once again, experimental results here along with those compared to the baseline models under our LFS setting reveal a *training data sparsity* issue that has to be addressed in any future multi-label zero-shot human action recognition study.

Furthermore, Table 7 shows the experimental results in conventional multi-label human action recognition, i.e., all the actions are known in learning. In this circumstance, only the IFS setting is applicable. Hence, we use the same IFS setting as described in Section 4.1.2 but, unlike what has been done for simulating a zero-shot scenario, do not reserve any actions. Also we use the same procedure as done for zero-shot learning to search for optimal hyper-parameters for five models and ours and repeat the experiments on the same data split as the IFS setting for three trials with different parameter initialization. As a result, we report the mean and standard deviation (std) of three-trial results yielded by different methods. It is evident from Table 7 that our model trained with either of two rank losses as well as their fusion outperform others in conventional multi-label recognition on both datasets. Without unseen classes, our model trained with the hinge rank loss generally performs slightly better than its counterpart trained with the RankNet loss. Once again, the fusion of results generated by those two models leads to the best performance. To see the degraded performance in a zero-shot scenario, we can compare the performance in the generalized ZSL evaluation scenario under the IFS setting, as shown in Table 6, to that reported in Table 7. By such a comparison, it is seen that the zero-shot performance of our model drops with a narrow margin (approximately less than 10% overall in terms of five different evaluation metrics). Given the fact that 10 out of 49 and 40 out of 157 human actions are reserved as unseen labels on Breakfast and Charades, respectively, this comparison on experimental results suggests that our proposed framework yields the promising performance for multi-label zero-shot human action recognition, which is close to the performance in multi-label human action recognition. Experimental results shown in Table 7 also suggest that other state-of-the-art methods behave similarly to ours in general. However, we also observe an unusual phenomenon from their performance; i.e., by a comparison to the generalized ZSL performance reported in Table 6, DSP yields slightly worse performance in multi-label recognition in terms of four of five evaluation metrics on Breakfast and so do Fast0Tag and

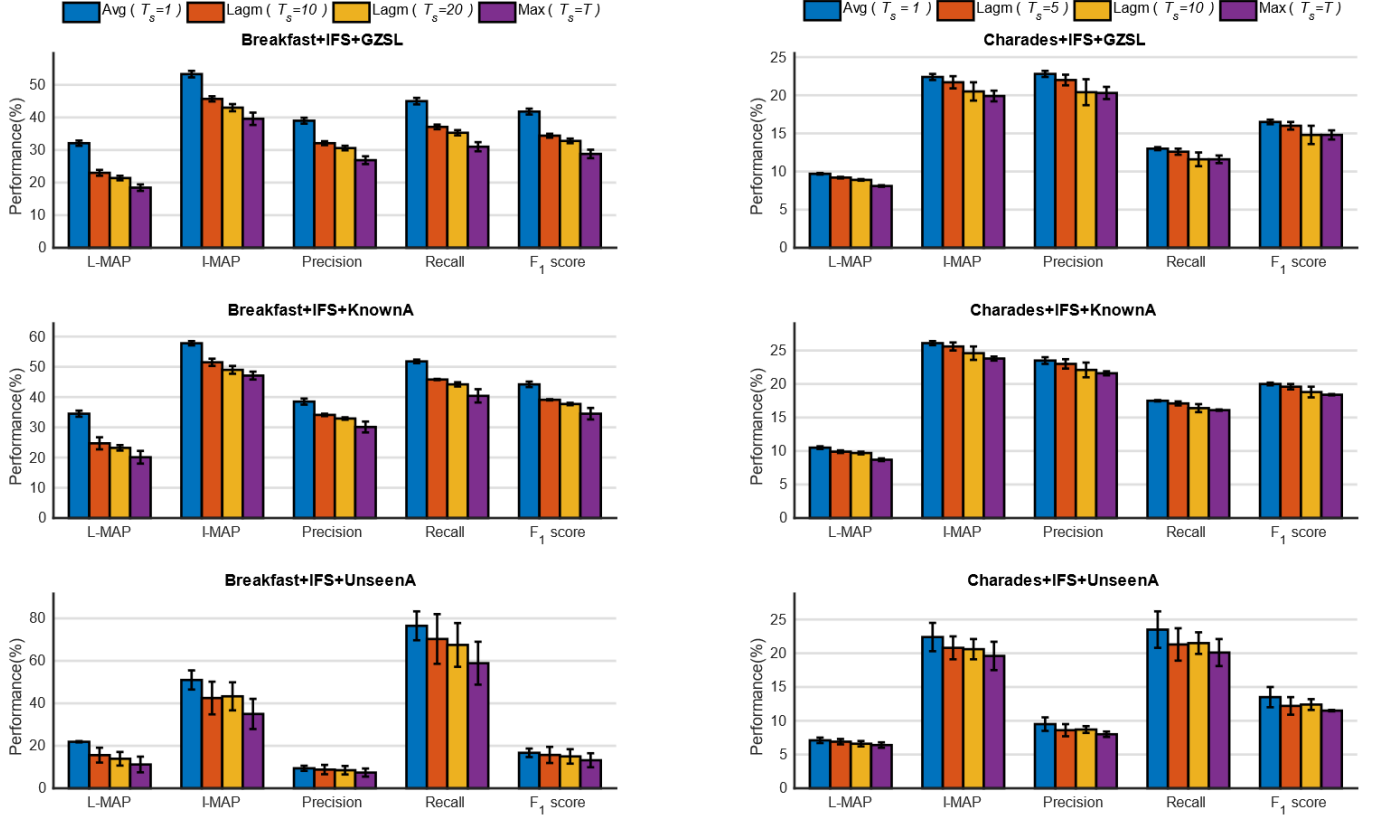


Figure 5: Performance of different pooling strategies used in our framework. Avg: average pooling; Lagm: local average global maximum pooling; Max: maximum pooling.

Fast0Tag+ in terms of L-MAP. By a closer look at the dataset and results in two experiments as well as our analysis, we find that at least two factors account for this unusual phenomenon: a) co-occurred labels associated with most of video clips on Breakfast are redundant in light of semantics, and b) the single collective semantic representation of co-occurred multiple labels used in DSP is insensitive to missing of few co-occurred labels due to the label information redundancy and the information loss resulting from the average operation in forming the single representation. Thus, we reckon that this phenomenon is rather specific to the nature of this dataset and the ZSL setting where there are only a small number of unseen labels.

5.4. Results on Pooling Strategy


We report the performance of three pooling strategies in terms of five evaluation criteria. It is evident from Figure 5 the average pooling always performs the best and the maximum pooling performs the worst regardless of evaluation criteria. In addition, the local average global maximum pooling performs better when T_s is set smaller. Such results imply that our framework interprets the visual information at a global level that tends to recognize actions appearing in a video clip rather than a local level that identifies the accurate boundaries between different actions. From our empirical study, it is observed that the average pooling takes into account all information in a video to yield the relatedness scores while the maximum pooling uses only the local information regarding an abrupt change in visual

domain but likely overlooks a large portion of useful information related to the nature of actions. Nevertheless, the maximum pooling might be beneficial for unsupervised action localization in the weak supervision setting, which is beyond the scope of this paper but worth studying in future.

In summary, our comparative study suggests that our proposed framework yields the favourable results and outperforms the existing state-of-the-art methods in general. The average pooling generally outperforms other alternatives in question. Also, our experimental results demonstrate challenges in multi-label ZSL via our novel LFS setting especially when training data are less correlated to test instances associated with unseen classes in both semantic and visual domains.


5.5. Visual Inspection

In general, visual inspection provides a manner that helps us understand the behaviour of a method intuitively. To gain an intuitive insight into the multi-label zero-shot human action recognition, we visualize a number of typical test video clips on Breakfast and the top-5 labels predicted by different state-of-the-art methods described in Section 4.6.2 and ours in terms of semantic relatedness scores. Our visual inspection mainly focuses on understanding of the behaviour of our model and issues arising from our work. As a result, Figures 6-9 illustrate several key frames to human actions in typical test video clips and the top-5 predicted labels by different methods, where a correctly predicted known label is highlighted with bold font




DSP	ConSE	COSTA	Fast0Tag	Fast0Tag+	Ours (RankNet)	Ours (Hinge)	Ours (Fusion)
<i>stir_egg</i> take_butter crack_egg pour_flour stir_dough	crack_egg stir_fry_egg <i>stir_egg</i> fry_egg stir_dough	crack_egg take_plate butter_pan stir_dough pour_milk	crack_egg fry_pancake add_salt_pepper put_egg2plate pour_milk	crack_egg butter_pan pour_milk stir_dough take_plate	crack_egg put_egg2plate fry_egg <i>stir_egg</i> fry_pancake	crack_egg take_plate butter_pan put_egg2plate stir_dough	crack_egg put_egg2plate take_plate butter_pan fry_pancake

Figure 6: A test video clip in the IFS setting and the top-5 labels predicted by different methods. Its ground-truth labels are **take_bowl**, **crack_egg**, **put_egg2plate**, **take_plate**, *stir_egg*, **pour_egg2pan**, **stir_fry_egg**, **add_salt_pepper**, **butter_pan**.



DSP	ConSE	COSTA	Fast0Tag	Fast0Tag+	Ours (RankNet)	Ours (Hinge)	Ours (Fusion)
pour_milk stir_egg pour_flour stir_milk pour_sugar	crack_egg pour_egg2pan fry_egg stir_fry_egg put_egg2plate	crack_egg put_egg2plate pour_milk add_salt_pepper fry_egg	crack_egg pour_milk fry_pancake put_egg2plate take_plate	pour_milk stir_milk spoon_powder add_teabag pour_water	pour_milk add_teabag spoon_powder pour_water stir_milk	pour_milk spoon_powder stir_milk spoon_sugar pour_water	pour_milk spoon_powder stir_milk add_teabag pour_water

Figure 7: A test video clip in the IFS setting and the top-5 labels predicted by different methods. Its ground-truth labels are **cut_orange**, **squeeze_orange**, **pour_juice**.



DSP	ConSE	COSTA	Fast0Tag	Fast0Tag+	Ours (RankNet)	Ours (Hinge)	Ours (Fusion)
take_butter stir_egg take_eggs crack_egg fry_egg	put_egg2plate crack_egg butter_pan pour_egg2pan fry_egg	crack_egg put_egg2plate butter_pan cut_bun add_salt_pepper	crack_egg add_salt_pepper put_egg2plate pour_milk fry_pancake	smear_butter cut_bun put_egg2plate add_salt_pepper crack_egg	crack_egg put_egg2plate fry_egg add_salt_pepper stir_fry_egg	crack_egg butter_pan take_plate put_egg2plate add_salt_pepper	crack_egg put_egg2plate take_palte butter_pan add_salt_pepper

Figure 8: A test video clip in the IFS setting and the top-5 labels predicted by different methods. Its ground-truth labels are **crack_egg**, **fry_egg**, **put_egg2plate**, **take_plate**, **add_salt_pepper**, **butter_pan**.



Data Split	DSP	ConSE	COSTA	Fast0Tag	Fast0Tag+	Ours(RankNet)	Ours(Hinge)	Ours (Fusion)
IFS	<i>stir_egg</i> crack_egg take_eggs take_butter pour_flour	pour_egg2pan pour_milk pour_flour pour_oil pour_sugar	put_egg2plate pour_milk crack_egg pour_oil add_salt_pepper	crack_egg fry_pancake put_egg2plate pour_milk add_salt_pepper	crack_egg pour_milk <i>stir_egg</i> put_egg2plate pour_oil	crack_egg put_egg2plate add_salt_pepper <i>stir_egg</i> fry_egg	crack_egg take_plate pour_oil put_egg2plate add_salt_pepper	crack_egg put_egg2plate take_plate add_salt_pepper Stir_fry_egg
LFS	squeeze_orange spoon_powder <i>stir_egg</i> stir_milk cut_bun	pour_sugar pour_flour pour_milk pour_juice pour_dough2pan	pour_sugar add_salt_pepper pour_milk spoon_powder put_bun_together	pour_milk spoon_powder fry_egg stir_milk pour_juice	pour_milk spoon_powder stir_milk cut_orange pour_juice	pour_juice squeeze_orange cut_orange take_galss take_plate	squeeze_orange cut_orange pour_juice take_glass take_plate	squeeze_orange cut_orange pour_juice take_plate take_galss

Figure 9: A test video clip appearing in the IFS and LFS settings and the top-5 labels predicted by different methods in two data split settings. Its ground truth labels are **take.bowl**, **crack_egg**, **put.egg2plate**, **take_plate**, *stir_egg*, **pour.egg2pan**, **stir_fry_egg**, **add_salt_pepper**, **butter_pan** in the IFS setting, and **take.bowl**, **crack_egg**, **put.egg2plate**, **take_plate**, *stir_egg*, **pour.egg2pan**, **stir_fry_egg**, **add_salt_pepper**, **butter_pan** in the LFS setting, respectively.

and a correctly predicted unseen label is marked with bold-italic font.

For the IFS setting, Figures 6-8 illustrate three typical results yielded by different methods. Figure 6 exemplifies the success of our model, where four out of the top-5 labels predicted by our model are the ground-truth actions and no other methods can match the performance of our model. This exemplified test instance suggests that the use of an LSTM layer in our visual model facilitates the recognition of distinctive actions in a video clip. Figure 7 shows a test instance where all the methods fail to have any ground-truth labels in their top-5 predicted labels. Our visual inspection on this test instance reveals that non-trivial objects pertaining to different actions are concentrated in a small region located in top-right of frames in this video clip. Thus, it is extremely difficult to capture the useful information in the visual domain, which poses a challenge to all the existing human action recognition techniques. Figure 6-9 reveal that our models trained with two rank losses yield different results for a test instance. Specially in Figure 9, three of the top-5 labels predicted by two models are in common, however, the fusion method described in Section 3.4 successfully predicts five ground truth labels. These instances vividly demonstrate the different aspects of two rank losses and the synergy achieved by their fusion. Besides, these test instances illustrated in Figures 6-9 also provide some insight regarding the behaviour of other state-of-the-art models used in our comparative study. For example, ConSE is more likely to yield the labels regarding frequently used words in a human action domain. For those test instances shown in Figures 6-8, at least four out of the top-5 labels predicted by ConSE are regarding different actions taken on “egg”. For the instance shown in Figure 9, all the top-5 labels predicted by ConSE are completely regarding “pour” actions commonly taken in kitchen. This limitation is due to the fact that ConSE uses a single collective semantic representation resulting from averaging the semantic representations of multiple co-occurred labels, which favors those frequently used

word vectors but diminishes the opportunity of finding out infrequently used word vectors in prediction.

Experimental results reported in Tables 4-6 suggest that all the models including ours generally perform worse under the LFS setting than under the IFS setting. On the one hand, the LFS setting results in a *training data sparsity* issue in contrast to the IFS setting. To see this issue, let us take the first split on Breakfast as an example. In this split shown in Table 2, there are 1,196 and 1,019 training examples in the IFS and the LFS setting, respectively. However, the number of training examples pertaining to specific known actions is significantly different in two split settings due to different data split protocols described in Section 4.1.2. For example, there are 330, 217, 254, 109 and 156 training examples with target labels, “crack.egg”, “put.egg2plate”, “take_plate”, “stir_fry_egg” and “add_salt_pepper”, respectively, in the IFS setting. In contrast, there are only 23, 23, 81, 23 and 20 examples with the same target labels, respectively, in the LFS setting. On the other hand, there is a major difference between those models and ours; i.e., our visual model employs a hidden layer of recurrent connection to capture temporal coherence underlying intrinsic visual features while those state-of-the-art models used in our comparative study do not have such a mechanism. It is well known that a learning model of a higher complexity or a larger capacity demands more informative training data. To this end, the training data sparsity issue affects the performance of our model more severely than other models; it is evident that the performance gain from the use of an LSTM layer in our visual model disappears due to a lack of sufficient training data required in training our visual model for capture temporal coherence.

To understand the difference between the IFS and the LFS settings and the training data sparsity issue intuitively, we illustrate the results yielded by the state-of-the-art methods and ours on a common instance appearing in test sets in two data split settings, as shown in Figure 9. It is evident that four out of

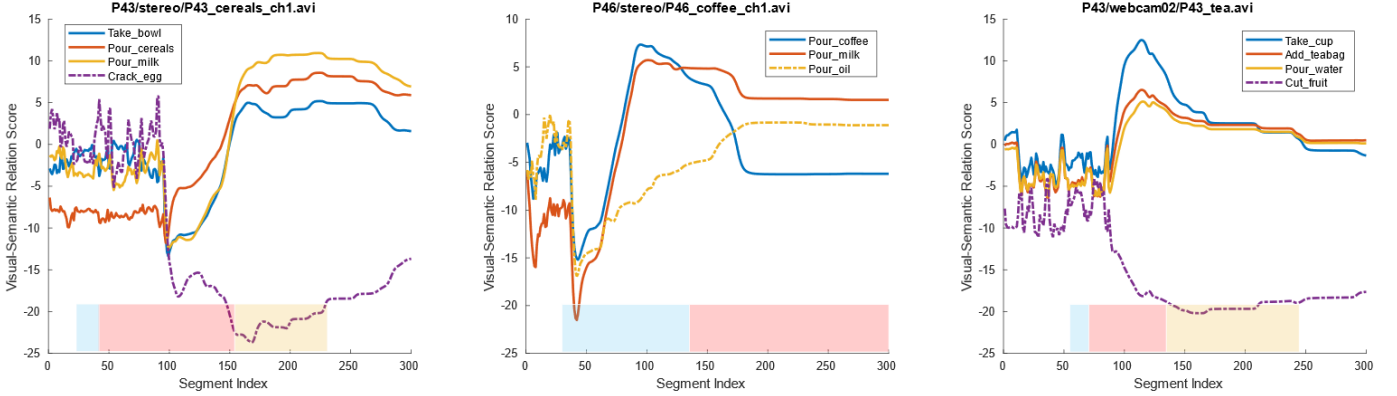


Figure 10: Visual-semantic relatedness scores of segments from three test video clips in the Breakfast dataset with the reference to the ground-truth location of actions. The horizontal bars (colored with blue, red and yellow) denote the ground-truth localization of different actions in a video clip. Their relatedness scores are depicted with solid lines of the corresponding colors while the dashed lines express the related scores of an illustrative negative action.

the top-5 action labels predicted by our model are the ground truth and all other models can predict some of ground-truth actions correctly under the IFS setting. In contrast, however, none of the models correctly predicts more than one ground-truth action for this exactly same test instance under the LFS setting. The visual inspection on this test instance clearly demonstrates the distinction between two data split settings; i.e., visual features associated with unseen actions are available in the IFS setting (an unrealistic scenario) but unavailable in the LFS setting (a realistic scenario), and the *training data sparsity* issue in the LFS setting, which poses a big challenge to all the existing multi-label ZSL methods including ours.

5.6. Correspondence between Visual and Semantic Embedding

To understand the behavior of our framework further, we explicitly investigate how our framework captures the correspondence between visual and semantic embedding by comparing the related scores against the ground truth of action boundaries further annotated manually. To this end, we take a closer look at the visual-semantic relatedness scores in terms of video segments corresponding to multiple actions in a video clip of a long duration.

In our experiment, we select three representative test video clips from the Breakfast dataset and exhibit the segment-level relatedness scores against the ground truth in Figure 10. The horizontal bars colored with blue, red and yellow in each plot express the ground-truth localization of actions in three video clips. Their relatedness scores are depicted with solid lines in the same colours, and the dashed lines denote the relatedness scores corresponding to some selected negative actions. It is observed from Figure 10 that the relatedness scores are usually below zero for all the actions at the beginning and with subtle changes within a period until reaching the point after which the relatedness scores of positive actions start increasing whilst the scores of negative actions get decreasing or stable. By comparing the moment when the significant change of relatedness scores occurs against the ground-truth localization of actions, we observe that the relatedness scores do not necessarily correspond to the accurate action location in a video clip. By a closer look at the left and middle plots in Figure 10, one can see that

the occurrence of a specific action would lead to a higher relatedness scores of this action than other positive actions not appearing within this visual segment during this period. However, this observation does not appear in the right plot. Hence, our experimental results suggest that our framework yields only reasonable relatedness scores once capturing sufficient visual information. The similar relatedness scores of multiple positive actions throughout the segment sequence indicate our framework learns the label co-occurrences well and hence recognize a set of actions collectively. However, it may fail in action localization which is beyond the scope of this paper.

5.7. Model Complexity

The architecture complexity of our learning model depends on the number of hidden layers, hidden units and their types as well as their connections used in neural networks to implement the visual and the semantic models for a data set.

In our current implementation, the number of parameters in the visual model varies from 4.9 to 24.7 millions, and the number of parameters in the semantic model varies from 0.15 to 0.77 millions under different hyper-parameter settings. Obviously, the semantic model has much fewer parameters compared with the visual model, suggesting that introducing a semantic model does not incur a much higher computational burden but leads to the performance gain. In general, our model often takes longer training time than other state-of-the-art learning models used in our comparative study due to the use of a LSTM layer to capture temporal coherence.

Practically, with a GTX1080Ti GPU, the averaging time spent for training our learning model is roughly 13 minutes on Charades (i.e., 40s per epoch multiplies approximately 20 epochs) and one hour on Breakfast due to a larger number of time steps ($T=300$). One limitation of our learning model is a large memory requirement for training. Recall that the visual representations have to be reserved for use in the training of semantic model, it is required to load one large matrix with a size of $n \times d_e \times T$ into memory. In our implementation, the amount of GPU memory used for training on Charades and Breakfast is 3.5 GB and 10.5 GB, respectively.

6. Concluding Remarks

In this paper, we have formulated human action recognition as a multi-label zero-shot learning problem and provide an effective solution by proposing a novel framework via joint latent ranking embedding learning. To carry out our framework, we employ a neural network of the heterogeneous architecture for visual embedding, where an LSTM layer is used to facilitate capturing temporal coherence information underlying different actions from weakly annotated video data. Also, we advocate the use of semantic embedding learning to facilitate bridging the semantic gap and effective knowledge transfer, which is implemented by a feed-forward neural network. All the above contributions have been thoroughly verified via our comparative study with various well-motivated settings. Experimental results on two benchmark multi-label human action datasets suggest that our proposed framework generally outperforms not only the baseline systems but also several state-of-the-art multi-label ZSL approaches in all the different test scenarios.

Although we have demonstrated favourable results on two benchmark datasets in comparison to state-of-the-art approaches, our observations on the performance of all the approaches used in our comparative study including ours suggest that the existing multi-label ZSL techniques are not ready for a real application; the instance-first split setting fails to simulate real multi-label zero-shot human action recognition scenarios while the performance becomes even worse under the label-first split setting that simulates a real scenario. Nevertheless, our experimental results including visual inspection provide the insightful information for improving our proposed framework. In our ongoing work, we would address issues arising from our experiments and observations with proper techniques. To address the *training data sparsity* issue revealed in our experiments, we would develop unsupervised learning algorithms to discover salient yet intrinsic visual features from unlabelled video clips and further incorporate proper temporal constraints into our rank loss functions to better capture temporal coherence. Moreover, we would consider diverse pooling strategies and introduce attention mechanisms to our model for improving implicit salient feature extraction and accurate localization of different yet complex actions involved in a video clip during the visual embedding learning. Also, we would employ alternative semantic representations developed by ourselves [63], which encode the semantic relatedness between action labels more accurately, in the semantic embedding learning to facilitate knowledge transfer.

While our framework is proposed especially for multi-label zero-shot human action recognition, we would highlight that it is directly applicable to multi-label human action recognition without modification as demonstrated in our experiments. Also, our framework is easy to adapt for tackling various multi-label ZSL problems in different domains. For example, we can apply our framework to miscellaneous multi-label zero-shot classification tasks on temporal or sequential data, e.g., acoustic event classification, straightforward as well as multi-label zero-shot learning tasks on static data, e.g., object recognition, by replacing a neural network of the heterogeneous architecture only

with a neural network of only feed-forward connections in the visual model. Thus, we are going to explore such extensions and applications in our future work.

Acknowledgement

The authors would like to thank the anonymous reviewers for their valuable comments that improve the presentation of this manuscript.

References

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2009, pp. 248–255.
- [2] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, S. Vijayanarasimhan, Youtube-8M: A large-scale video classification benchmark, arXiv preprint arXiv:1609.08675.
- [3] M. L. Zhang, Z. H. Zhou, A review on multi-label learning algorithms, IEEE Transactions on Knowledge and Data Engineering 26 (8) (2014) 1819–1837.
- [4] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, A. Gupta, Hollywood in homes: Crowdsourcing data collection for activity understanding, in: European Conference on Computer Vision (ECCV), Springer, 2016, pp. 510–526.
- [5] Y. Zhang, R. Acharyya, J. Liu, B. Gong, Infinite-label learning with semantic output codes, arXiv preprint arXiv:1608.06608.
- [6] Y. Zhang, B. Gong, M. Shah, Fast zero-shot image tagging, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2016, pp. 5985–5994.
- [7] M. Guillaumin, T. Mensink, J. Verbeek, C. Schmid, Tagprop: Discriminative metric learning in nearest neighbor models for image auto-annotation, in: IEEE International Conference on Computer Vision (ICCV), IEEE, 2009, pp. 309–316.
- [8] J. Wang, Y. Yang, J. Mao, Z. Huang, C. Huang, W. Xu, CNN-RNN: A unified framework for multi-label image classification, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 2285–2294.
- [9] J. Nam, J. Kim, E. L. Mencia, I. Gurevych, J. Fürnkranz, Large-scale multi-label text classification – revisiting neural networks, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD), Springer, 2014, pp. 437–452.
- [10] Z.-H. Zhou, M.-L. Zhang, Multi-instance multi-label learning with application to scene classification, Advances in neural information processing systems (NIPS) 19 (2007) 1609.
- [11] F. Gu, M. Sridhar, A. Cohn, D. Hogg, F. Flórez-Revue, D. Monekosso, P. Remagnino, Weakly supervised activity analysis with spatio-temporal localisation, Neurocomputing 216 (2016) 778–789.
- [12] Y. Wei, W. Xia, M. Lin, J. Huang, B. Ni, J. Dong, Y. Zhao, S. Yan, HCP: A flexible CNN framework for multi-label image classification, IEEE Transactions on Pattern Analysis and Machine Intelligence 38 (9) (2016) 1901–1907.
- [13] P. Tang, X. Wang, Z. Huang, X. Bai, W. Liu, Deep patch learning for weakly supervised object classification and discovery, Pattern Recognition 71 (3) (2017) 446–459.
- [14] R. Cabral, F. De la Torre, J. P. Costeira, A. Bernardino, Matrix completion for weakly-supervised multi-label image classification, IEEE Transactions on Pattern Analysis and Machine Intelligence 37 (1) (2015) 121–135.
- [15] Q. Wang, K. Chen, Zero-shot visual recognition via bidirectional latent embedding, International Journal of Computer Vision 124 (3) (2017) 356–383.
- [16] C. H. Lampert, H. Nickisch, S. Harmeling, Attribute-based classification for zero-shot visual object categorization, IEEE Transactions on Pattern Analysis and Machine Intelligence 36 (3) (2014) 453–465.
- [17] Y. Fu, T. Huang, Manifold and subspace learning for pattern recognition, Pattern Recognition and Machine Vision 6 (2010) 215.

- [18] Z. Akata, S. Reed, D. Walter, H. Lee, B. Schiele, Evaluation of output embeddings for fine-grained image classification, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 2927–2936.
- [19] Y. Xian, Z. Akata, G. Sharma, Q. Nguyen, M. Hein, B. Schiele, Latent embeddings for zero-shot classification, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [20] Y. Yu, Z. Ji, J. Guo, Y. Pang, Transductive zero-shot learning with adaptive structural embedding, IEEE transactions on neural networks and learning systems 29 (9) (2018) 4116–4127.
- [21] S. Changpinyo, W.-L. Chao, B. Gong, F. Sha, Synthesized classifiers for zero-shot learning, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [22] T. Mensink, E. Gavves, C. G. Snoek, Costa: Co-occurrence statistics for zero-shot classification, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 2441–2448.
- [23] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al., Devise: A deep visual-semantic embedding model, in: Advances in neural information processing systems (NIPS), 2013, pp. 2121–2129.
- [24] J. Lei Ba, K. Swersky, S. Fidler, et al., Predicting deep zero-shot convolutional neural networks using textual descriptions, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2015, pp. 4247–4255.
- [25] S. Changpinyo, W.-L. Chao, F. Sha, Predicting visual exemplars of unseen classes for zero-shot learning, in: IEEE International Conference on Computer Vision, 2017.
- [26] L. Zhang, T. Xiang, S. Gong, Learning a deep embedding model for zero-shot learning, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- [27] Z. Zhang, V. Saligrama, Zero-shot learning via semantic similarity embedding, in: IEEE International Conference on Computer Vision (ICCV), 2015, pp. 4166–4174.
- [28] Z. Zhang, V. Saligrama, Zero-shot learning via joint latent similarity embedding, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 6034–6042.
- [29] Y. Yu, Z. Ji, J. Guo, Z. Zhang, Zero-shot learning via latent space encoding, IEEE transactions on cybernetics (99) (2018) 1–12.
- [30] Z. Ren, H. Jin, Z. Lin, C. Fang, A. Yuille, Multi-instance visual-semantic embedding, arXiv preprint arXiv:1512.06963.
- [31] U. Sandouk, K. Chen, Multi-label zero-shot learning via concept embedding, arXiv preprint arXiv:1606.00282.
- [32] J. Nam, E. L. Mencía, H. J. Kim, J. Fürnkranz, Predicting unseen labels using label hierarchies in large-scale multi-label learning, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD), Springer, 2015, pp. 102–118.
- [33] Y. Fu, Y. Yang, T. Hospedales, T. Xiang, S. Gong, Transductive multi-label zero-shot learning, in: Proceedings of British Machine Vision Conference (BMVC), 2014.
- [34] J. Weston, S. Bengio, N. Usunier, Large scale image annotation: Learning to rank with joint word-image embeddings, Machine learning 81 (1) (2010) 21–35.
- [35] C. Xu, D. Tao, C. Xu, A survey on multi-view learning (2013). [arXiv: 1304.5634](https://arxiv.org/abs/1304.5634).
- [36] Y. Gong, Q. Ke, M. Isard, S. Lazebnik, A multi-view embedding space for modeling internet images, tags, and their semantics, International Journal of Computer Vision 106 (2) (2014) 210–233.
- [37] A. Karpathy, L. Fei-Fei, Deep visual-semantic alignments for generating image descriptions, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 3128–3137.
- [38] Y. Xian, B. Schiele, Z. Akata, Zero-shot learning - the good, the bad and the ugly, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- [39] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Advances in neural information processing systems (NIPS), 2013, pp. 3111–3119.
- [40] U. Sandouk, K. Chen, Learning contextualized semantics from co-occurring terms via a siamese architecture, Neural Networks 76 (2016) 65–96.
- [41] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, G. Toderici, Beyond short snippets: Deep networks for video classification, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 4694–4702.
- [42] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 2625–2634.
- [43] C.-Y. Ma, M.-H. Chen, Z. Kira, G. AlRegib, TS-LSTM and temporal-inception: Exploiting spatiotemporal dynamics for activity recognition, arXiv preprint arXiv:1703.10667.
- [44] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Computation 9 (8) (1997) 1735–1780.
- [45] V. Nair, G. Hinton, Rectified linear units improve restricted boltzmann machines, in: International Conference on Machine Learning (ICML), 2010.
- [46] A. Graves, Generating sequences with recurrent neural networks, arXiv preprint arXiv:1308.0850.
- [47] M. Lapin, M. Hein, B. Schiele, Analysis and optimization of loss functions for multiclass, top-k, and multilabel classification, IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [48] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, G. Hullender, Learning to rank using gradient descent, in: Proceedings of the 22nd international conference on Machine learning (ICML), ACM, 2005, pp. 89–96.
- [49] R. Herbrich, K. Obermayer, T. Graepel, Large margin rank boundaries for ordinal regression, in: Advances in Large Margin Classifiers, MIT Press, Cambridge MA, 1999, pp. 115–132.
- [50] K. Kavukcuoglu, M. Ranzato, Y. LeCun, Fast inference in sparse coding algorithms with applications to object recognition, arXiv preprint arXiv:1010.3467.
- [51] Y.-G. Jiang, Z. Wu, J. Wang, X. Xue, S.-F. Chang, Exploiting feature and class relationships in video categorization with regularized deep neural networks, IEEE Transactions on Pattern Analysis and Machine Intelligence (2017) (in press).
- [52] H. Kuehne, A. Arslan, T. Serre, The language of actions: Recovering the syntax and semantics of goal-directed human activities, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 780–787.
- [53] D. Tran, L. Bourdev, R. Fergus, L. Torresani, M. Paluri, Learning spatiotemporal features with 3d convolutional networks, in: IEEE International Conference on Computer Vision (ICCV), IEEE, 2015, pp. 4489–4497.
- [54] H. Wang, C. Schmid, Action recognition with improved trajectories, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2013, pp. 3551–3558.
- [55] F. Chollet, Keras, <https://github.com/fchollet/keras> (2015).
- [56] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Proceedings of the 3rd International Conference on Learning Representations (ICLR), 2014.
- [57] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, Journal of Machine Learning Research 15 (1) (2014) 1929–1958.
- [58] M. S. Sorower, A literature survey on algorithms for multi-label learning, Technical Report, Oregon State University, Corvallis, U.S.A.
- [59] X. Li, T. Uricchio, L. Ballan, M. Bertini, C. G. Snoek, A. D. Bimbo, Socializing the semantic gap: A comparative survey on image tag assignment, refinement, and retrieval, ACM Computing Surveys 49 (1) (2016) 14.
- [60] Y. Gong, Y. Jia, T. Leung, A. Toshev, S. Ioffe, Deep convolutional ranking for multilabel image annotation, arXiv preprint arXiv:1312.4894.
- [61] C. D. Manning, P. Raghavan, H. Schtze, An introduction to inf. retrieval, in: An Introduction to Inf. Retrieval, Cambridge University Press, 2009.
- [62] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, J. Dean, Zero-shot learning by convex combination of semantic embeddings, in: International Conference on Learning Representations (ICLR), 2014.
- [63] Q. Wang, K. Chen, Alternative semantic representations for zero-shot human action recognition, in: Proceedings of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery (ECML-PKDD), 2017.