

Structured Label Inference for Visual Understanding

by

Nelson Nauata

B.Eng., University of Campinas, 2016

Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
Department of Computing Science
Faculty of Applied Sciences

© Nelson Nauata 2018
SIMON FRASER UNIVERSITY
Spring 2018

Copyright in this work rests with the author. Please ensure that any reproduction or re-use is done in accordance with the relevant national copyright legislation.

Approval

Name: Nelson Nauata

Degree: Master of Science (Computer Science)

Title: Structured Label Inference for Visual Understanding

Examining Committee: **Chair:** Dr. KangKang Yin
Associate Professor

Dr. Greg Mori
Senior Supervisor
Professor

Dr. Richard Zhang
Co-Supervisor
Professor

Dr. Yasutaka Furukawa
Internal Examiner
Assistant Professor

Date Defended: April 18th, 2018

Abstract

Visual data such as images and videos contain a rich source of structured semantic labels as well as a wide range of interacting components. Visual content could be assigned with fine-grained labels describing major components, coarse-grained labels depicting high level abstractions, or a set of labels revealing attributes. Such categorization over different, interacting layers of labels evinces the potential for a graph-based encoding of label information. In this thesis, we exploit this rich structure for performing graph-based inference in label space for a number of tasks: multi-label image and video classification and action detection in untrimmed videos. We consider the use of the Bidirectional Inference Neural Network (BINN) and Structured Inference Neural Network (SINN) for performing graph-based inference in label space and propose a Long Short-Term Memory (LSTM) based extension for exploiting activity progression on untrimmed videos. The methods were evaluated on (i) the Animal with Attributes (AwA), Scene Understanding (SUN) and NUS-WIDE datasets for multi-label image classification, (ii) the first two releases of the YouTube-8M large scale dataset for multi-label video classification, and (iii) the THUMOS'14 and MultiTHUMOS video datasets for action detection. Our results demonstrate the effectiveness of structured label inference in these challenging tasks, achieving significant improvements against baselines.

Keywords: Computer vision, multi-label classification, image classification, video recognition, action detection, structured inference.

Acknowledgements

I would like to thank my colleagues Hexiang Hu, Dr. Guang-Tong Zhou and Dr. Zicheng Liao for helping me completing and peer reviewing this work. Furthermore, I would like to thank my supervisor, Dr. Greg Mori, for patient guidance and encouragement throughout my time as his student, which were crucial for the completion of this project. Many thanks to my examining committee for reading the thesis. Special thanks to my wife, Fernanda, for supporting me throughout the entire process.

Table of Contents

Approval	ii
Abstract	iii
Acknowledgements	iv
Table of Contents	v
List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Contribution	4
2 Related Work	5
2.1 Label Relations and External Structured Concepts	5
2.2 Multi-task Joint Learning	5
2.3 Deep Visual Recognition	6
2.4 Video Classification	6
2.5 Action Detection	6
2.6 Recurrent Temporal Modeling	7
3 Structured Inference Framework	8
3.1 Bidirectional Inference Neural Network (BINN)	9
3.2 Structured Inference Neural Network (SINN)	10
3.3 Modeling Cross-temporal Relations	12
3.4 Prediction Framework	14
4 Multi-label Classification on Images	15
4.1 Evaluation Benchmarks	16
4.1.1 Animal with Attributes (AwA)	16
4.1.2 Scene Understanding (SUN)	16
4.1.3 Real-World Web Images (NUS-WIDE)	16

4.2	Evaluation Metrics	16
4.2.1	Multiclass Accuracy (MC Acc)	16
4.2.2	Intersection over Union Accuracy (IoU Acc)	17
4.2.3	Average Precision (AP)	17
4.2.4	Per-label Mean Average Precision (mAP_L)	17
4.2.5	Per Image Mean Average Precision (mAP_I)	17
4.3	Experimental Setup	18
4.3.1	AwA: Layered Prediction with Label Relations	18
4.3.2	NUS-WIDE: Multi-label Classification with Partial Human Labels of Tags and Groups	18
4.3.3	SUN397: Improving Scene Recognition with and without partially Ob- served Labels	19
4.4	Implementation Details	19
4.5	Experimental Results	19
5	Multi-label Classification on Videos	23
5.1	Evaluation Benchmarks	24
5.1.1	YouTube-8M (YT-8M)	24
5.1.2	YouTube-8M V2 (YT-8M V2)	24
5.2	Evaluation Metrics	24
5.2.1	Hit at Top K (Hit@k)	24
5.2.2	Precision at Equal Recall Rate (PERR)	25
5.2.3	Global Average Precision (gAP)	25
5.2.4	Per Video Mean Average Precision (mAP_V)	25
5.3	Experimental Setup	26
5.4	Implementation Details	26
5.5	Experimental Results	26
6	Dense Action Detection	28
6.1	Evaluation Benchmarks	29
6.1.1	THUMOS'14	29
6.1.2	MultiTHUMOS	29
6.2	Evaluation Metrics	29
6.3	Experimental Setup	29
6.4	Implementation Details	30
6.5	Experimental Results	30
7	Conclusion	35
	Bibliography	36

List of Tables

Table 4.1	Layered label prediction results on the AwA dataset.	20
Table 4.2	Layered label prediction results on the SUN397 dataset.	21
Table 4.3	Recognition results on the 397 fine-grained scene categories. Note that the last two compared methods make use of partially observed labels from the top (coarsest) scene layer, i.e. <i>indoor</i> , <i>outdoor man-made</i> , and <i>outdoor natural</i>	21
Table 4.4	Results on NUS-WIDE for from [15]. We also included the results for per-label recall and precision (R_L and P_L) and per-image (R_I and P_I), considering the top 3 labels for each image.	22
Table 5.1	YouTube-8M (YT-8M) results for mAP_V , Precision at Equal Recall Rate (PERR), Hit at 1 and gAP on the validation set.	27
Table 5.2	YouTube-8M V2 (YT-8M V2) results for mAP_V , Precision at Equal Recall Rate (PERR), Hit at 1 and gAP on the validation set.	27
Table 6.1	THUMOS and MultiTHUMOS results for mAP_L	31

List of Figures

Figure 1.1	This image example has visual concepts at various levels, from sports field at a high level to baseball and person at lower levels. Our model leverages label relations and jointly predicts layered visual labels from an image using a structured inference neural network. In the graph, colored nodes correspond to the labels associated with the image, and red edges encode label relations.	2
Figure 3.1	The label prediction pipeline. Given an input image, we extract CNN features at the last fully-connected layer as activation (in blue box) at different visual concept layers. We then propagate the activation information in a label (concept) relation graph through our structured inference neural network (in red box). The final label prediction (in green box) is made from the output activations (in yellow box) obtained after the inference process.	8
Figure 3.2	An example showing the model parameters V_p and V_n between the animal layer and the attribute layer. Green edges in the graph represent positive correlation, and red edges represent negative correlation.	11
Figure 3.3	The label prediction pipeline with partial observation. The pipeline is similar to Figure 3.1 except that we now have a partial observation that this image is outdoor man-made. The SINN is able to take the observed label into consideration and improve the label predictions in the other concept layers.	12
Figure 4.1	The inference model is applied directly on a feature vector extracted by a CNN for performing multi-label image classification.	15
Figure 5.1	The video-level representation $\bar{\mathbf{x}}$ is obtained by pooling per-frame feature vectors \mathbf{x}_t (using a CNN) and fed to the inference model for performing multi-label video classification.	23

Figure 6.1	Snapshot at time step t for the action detection model, where per-frame representations \mathbf{x}_t were extracted using a CNN, fed to the inference model and to m concept-layer specific LSTM units at each time step t . The concept layer outputs $\{\mathbf{a}_t^\ell\}_{\ell=1}^m$ are combined (aggregation step denoted by σ and formulated in Eq. (3.19)) with the corresponding LSTM hidden states $\{\mathbf{h}_t^\ell\}_{\ell=1}^m$ for obtaining the final predictions $\{\mathbf{y}_t^\ell\}_{\ell=1}^m$. The output activations (\mathbf{y}_t^ℓ) for the concept layer ℓ are interpreted as confidence scores for a given concept being assigned to the frame at time step t	28
Figure 6.2	Per-class APs for Single-frame SINN versus Single-frame CNN (VGG).	32
Figure 6.3	Per-class APs for LSTM versus Single-frame CNN (VGG).	32
Figure 6.4	Per-class APs for siLSTM versus LSTM.	33
Figure 6.5	Per-class APs for siLSTM versus Single-frame CNN (VGG).	33
Figure 6.6	Timeline comparison for single-frame models (CNN and SINN) and sequential models (LSTM and siLSTM) in multi-label action detection for four different videos. The vertical axis represents the ground-truth labels appearing at least once in the target video and the horizontal axis corresponds to the duration of the target video in minutes. (Best view in color)	34

Chapter 1

Introduction

Visual content is a rich source of high-dimensional structured data, with a wide range of interacting components at varying levels of abstractions. With the proliferation of large-scale image [41, 3, 30, 54] and video [1, 20, 57] datasets, advances in visual understanding were facilitated for the exploration and enhancement of intelligent reasoning techniques for modelling structured concepts. In this thesis, we exploit these rich structures for modelling concept interactions in a number of different tasks and levels of complexity: multi-label image classification, multi-label video classification and action detection in untrimmed videos.

Standard image classification is a fundamental problem in computer vision – assigning category labels to images. It can serve as a building block for many different computer vision tasks including object detection, visual segmentation, scene parsing and concept localization. Successful deep learning approaches [28, 42, 44, 47] typically assume labels to be semantically independent and adapt either a multi-class or binary classifier to target labels. In recent work [4, 6], deep learning methods that take advantage of label relations have been proposed to improve classification performance. However, in realistic settings, these label relationships could form a complicated graph structure. Take Figure 1.1 as an example. Various levels of interpretation could be formed to represent such visual content. This image of a *baseball* scene could be described as an *outdoor* image at coarse level, or with a more concrete concept such as *sports field*, or with even more fine-grained labels such as *batter's box* and objects such as *grass*, *bat*, *person*.

Models that incorporate semantic label relationships could be utilized to generate better classification results. The desiderata for these models include the ability to model label-label relations such as positive or negative correlation, respect multiple concept layers obtainable from sources such as WordNet, and to handle partially observed label data - given a subset of accurate labels for this image, infer the remaining missing labels.

In our previous work [15], we developed a structured inference neural network that permits modeling complex relations between labels, ranging from hierarchical to within-layer dependencies. We achieve this by defining a network in which a node is activated if its corresponding label is present in an image. We introduce stacked layers among these label

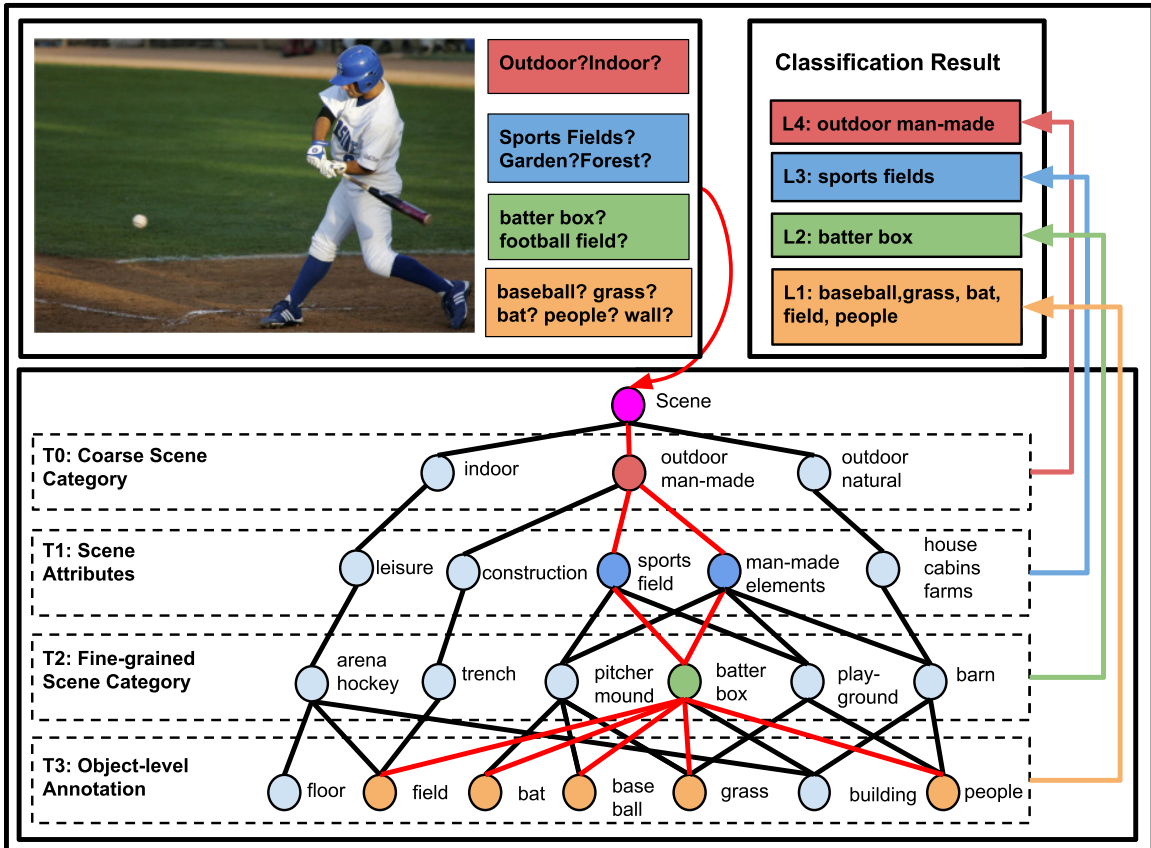


Figure 1.1: This image example has visual concepts at various levels, from sports field at a high level to baseball and person at lower levels. Our model leverages label relations and jointly predicts layered visual labels from an image using a structured inference neural network. In the graph, colored nodes correspond to the labels associated with the image, and red edges encode label relations.

nodes. These encode layer-wise connectivity among label classification scores, representing dependencies from top-level coarse labels to bottom-level fine-grained labels. Activations are propagated bidirectionally and asynchronously on the label relation graph, passing information about the labels within or across concept layers to refine the labeling for the entire image.

The similarity between multi-label classification on images and videos suggests exploitation of structured data to be beneficial for both tasks. As demonstrated in [38], our method extends beyond its applicability on images and is robust to higher dimensional structured data such as videos. A more challenging problem than multi-label video classification consists of handling a sequential input of frames and inferring the corresponding sequence of dense annotations. Our exploration of this setting for the task of dense action detection [57] is two-fold: performing static frame structured inference and spatio-temporal structured inference.

The static frame inference can be reduced to a standard image classification problem. However, spatio-temporal structured inference requires modelling cross-temporal relationships between labels. A natural way of extending [38] to support this feature is allowing communication between concept layers in the hierarchical structure for forward propagating learned label correlations and exploring labels' progression on untrimmed videos. We achieve this by enriching the hidden state of Long Short-Term Memory (LSTM) [14] units with extracted information from our structured inference method along the frame sequences.

The contribution of this thesis is binding together the proposed model presented in [15] for performing hierarchical inference on image datasets (AwA [30] and SUN397 [54]) and its video extension previously implemented in [38]. In addition, we include novel results for the two recent releases of Youtube-8M [1], including partially observed labels and propose a temporal extension for the bidirectional and structured inference models, demonstrating that adding cross-temporal information in label space (i.e. propagation across concept layers) provides superior performance against a traditional technique for incorporating temporal dependencies (i.e. LSTM). The validation of the proposed models was carried out on THUMOS [20] and MultiTHUMOS [57] for the task of action detection.

This thesis is organized as follows. Chapter 2 presents prior knowledge, covering previous related work. In Chapter 3, we bind together the bidirectional and structured inference models formulated in [15] with its video extension presented in [38] and derive the proposed formulation for temporal extension of our bidirectional and structured inference methods. Chapters 4 and 5 present the work done on multi-label image and video classification respectively. Chapter 6 describes the work done on the action detection task. We conclude in Chapter 7 with a brief discussion.

1.1 Contribution

The thesis is based on a paper co-authored by Hexiang Hu, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, and Greg Mori. Learning structured inference neural networks with label relations. In *CVPR*, 2016.

My contributions to this work were to develop the extensions to video and run all experiments on the two releases of Youtube-8M, including experiments for baselines, BINN, SINN and the partially observed settings. In addition, I was responsible for developing the temporal extensions for action detection denoted biLSTM and siLSTM, plus validating these models on THUMOS'14 and MultiTHUMOS, which included reporting quantitative and qualitative results.

Chapter 2

Related Work

Structured labeling information has been incorporated in numerous frameworks throughout previous literature. In this section, we review some prominent related works and the algorithms implemented for the various tasks studied in this thesis.

2.1 Label Relations and External Structured Concepts

The incorporation of label relations is often explored by modelling graphical structures on the training data (e.g. [5]) or constraining a loss function for jointly predicting structured labels [49, 50]. This work investigates the ability to model label relations such as positive and negative correlations, building on top of the bidirectional and structured inference neural network previously proposed in [15].

When external structured concepts are available it is advantageous to incorporate them to conduct traditional supervised approaches. For instance, Grauman *et al.* [10] and Hwang *et al.* [16] exploit the WordNet taxonomy for learning a discriminative tree of metrics of visual representations hierarchically structured. Johnson *et al.* [21] and McAuley and Leskovec [35] exploited social-network metadata, taking into account its interdependencies for applying structured learning and enhancing image classification against methods relying solely on image content.

2.2 Multi-task Joint Learning

The intuition behind structured label prediction is closely related to multi-task learning, with the distinction that multiple correlated tasks are jointly estimated. Common jointly modeled tasks include segmentation and detection [29, 53], segmentation and pose estimation [26], or segmentation and object classification [31]. An emerging topic of joint learning lies in image understanding and text generation by leveraging intra-modal correspondences between visual data and human language [27, 22].

The method presented in [15] might also be extended to multi-task learning. A natural way of achieving this is considering each concept layer to be a different task, whose labels do

not necessarily contain a layered structure. Thus, the existing multi-task learning methods might also benefit from prior knowledge of intra-task relations.

2.3 Deep Visual Recognition

The use of deep learning for applications in videos has advanced in lockstep with the field’s success in images. A traditional deep learning framework for extracting spatial descriptions comprises using convolutional neural networks (CNNs) such as AlexNet [28], VGG-16 [44] or InceptionV3 [48], pre-trained on a diverse and large image collection (e.g. ImageNet [41]). Recently, deep residual networks [12] built on these successes, and pushed the image classification performance of deep CNNs to impressive performance levels on benchmark datasets. In this thesis, we leverage those recent progresses for structured label inference.

2.4 Video Classification

Traditional approaches for video recognition [52, 51] build classifiers on top of hand-crafted features such as Histograms of Oriented Gradients (HOG), Histograms of Optical Flow (HOF) and Motion Boundary Histograms (MBH). On the contrary, common deep learning strategies for improving recognition consist of automatically learning descriptors that capture discriminative appearance and motion features via implementation of spatio-temporal networks [23, 18, 2] or exploitation of temporal dependencies using recurrent neural networks [39, 7]. Instead our approach intends to leverage structured label information for performing stronger video-level label predictions. For this reason, we adopt the simple frame-level feature aggregation from [1], which consists of average pooling spatial features across untrimmed YouTube videos [1] for performing video classification.

2.5 Action Detection

More challenging than performing video-level recognition is localizing concepts at a frame-level. In this thesis, we address the problem of dense multi-label action detection. We utilize a recent dataset [57], which requires inferring multiple frame-level labels across videos. A classic approach before deep learning [24] tackles the problem of detection by matching volumetric representations of events against oversegmented videos. Aligned more closely to our approach are [45, 37, 34], which analyze label progression, operating on short clips or single frames for inferring actions for each time step using recurrent neural networks. In contrast, we consider label progression over a structured collection of labels.

2.6 Recurrent Temporal Modeling

Early attempts of modelling time-sequential images in computer vision rely on hidden Markov models for obtaining representations in order to recognize action classes considering evolution of actions or even multiple view-points [56, 33, 43]. A stronger presence of deep learning in recent years along with its eminent success in the field, shifted attention to recurrent neural network (RNN) approaches for modelling time dependent correlations in sequential data. In particular for frame sequences, LSTM based models have been extensively explored recently [39, 7, 57] due to their capacity of efficiently transmitting information across time. In this thesis, we leverage such successes of temporal modelling techniques for extending our structured label inference framework across time.

Chapter 3

Structured Inference Framework

Our model jointly classifies images or videos in a layered label space with external label relations. The goal is to leverage the label relations to improve inference over the layered visual concepts and extend our approach for modelling temporal dependencies for action detection.

We build our model on top of a state-of-the-art deep learning platform: given a collection of training images \mathcal{I} (i.e. $\{I^i\}_{i=1}^{|\mathcal{I}|}$) or a collection of videos \mathcal{V} (i.e. $\{V^j\}_{j=1}^{|\mathcal{V}|}$), where the j -th video is composed by frames $\{F_k^j\}_{k=1}^{|V^j|}$, we first reduce spatial dimensions by extracting CNN features as visual activations. For example, a feature vector \mathbf{x}^i for the i -th image in the collection is obtained by $CNN(I^i)$, or a feature vector \mathbf{x}_k^j for the j -th video and k -th frame is obtained by applying $CNN(F_k^j)$.




Figure 3.1: The label prediction pipeline. Given an input image, we extract CNN features at the last fully-connected layer as activation (in blue box) at different visual concept layers. We then propagate the activation information in a label (concept) relation graph through our structured inference neural network (in red box). The final label prediction (in green box) is made from the output activations (in yellow box) obtained after the inference process.

In the particular case of inputting single images or frames (i.e. I^i or F_k^j), spatial representations \mathbf{x}^i or \mathbf{x}_k^j , can be fed directly to our model. For an entire video $\{F_k^j\}_{k=1}^{|V^j|}$, we deal with inputs in either of two ways: providing a summarized feature vector $\bar{\mathbf{x}}^j$ for the entire video or a sequence of feature vectors $\{\mathbf{x}_k^j\}_{k=1}^{|V^j|}$ extracted from the corresponding input frames. For video classification, we apply the former case by including an intermediate aver-

aging pooling across the entire video [1] and the latter case serves as input to our temporal extension model.

In Chapter 3.1 and Chapter 3.2, we formulate our models for tagging with multiple labels \mathbf{y}^i a single image I^i , or with multiple video-level labels \mathbf{y}^j an entire video V^j . In Chapter 3.3, we describe our proposed temporal extension method for outputting multiple per-frame actions as continuous confidence scores $\{\mathbf{y}_k^j\}_{k=1}^{|V_j|}$ in a video V_j , using a temporal sequence of spatial feature vectors $\{\mathbf{x}_k^j\}_{k=1}^{|V_j|}$ as input. In Chapter 3.4, we present our approach for performing structured predictions and including partial observations in our models.

3.1 Bidirectional Inference Neural Network (BINN)

Our model is inspired by the recent success of RNNs [11, 32], which make use of dynamic sequential information in learning. RNNs are called recurrent models because they perform the same computation for every time step, with the input dependent on the current inputs and previous outputs. Applying a similar idea to our layered label prediction problem: we consider each concept layer as an individual time step, and model the label relations within and across concept layers in the recurrent learning framework.

The learning of structured label relations is seen as a hierarchical distribution of labels in our architecture, where each level is defined as a concept layer and represents the degree of granularity for the label space. In the implementation for the BINN, the labels are separated into a set \mathcal{M} of concept layers, with varying granularity, totalling m layers. For example, a coarse-grained label for a scene could be *outdoors*, whereas a fine-grained label would be *tree*.

From now on, we denote a single input feature vector $\mathbf{x} \in \mathbb{R}^D$ with ground-truth labels as a binary vector \mathbf{t}^ℓ at concept layer $\ell \in \mathcal{M}$, where $t_k^\ell = 1$, if the k -th concept is present and $t_k^\ell = 0$, otherwise. This feature vector will be treated as input separately for each concept layer. The concept layers are represented as activations obtained by performing inference in the graph – for each concept layer ℓ , we have an activation vector $\mathbf{a}^\ell \in \mathbb{R}^{n_\ell}$ associated with the labels at concept layer ℓ , where n_ℓ is the number of labels at concept layer ℓ . In order to perform inference, the dimension of the input \mathbf{x} should be regressed to the label space. Thus, the input \mathbf{x}^ℓ for each concept layer ℓ and an input feature vector \mathbf{x} is given by:

$$\mathbf{x}^\ell = W^\ell \cdot \mathbf{x} + \mathbf{b}^\ell, \quad (3.1)$$

where $W^\ell \in \mathbb{R}^{n_\ell \times D}$ and $\mathbf{b}^\ell \in \mathbb{R}^{n_\ell \times 1}$ are learnable parameters.

The bidirectional message passing consists of two steps: a top-down inference and a bottom-up inference. The former captures inter-layer and intra-layer label relations in the top-down direction computing intermediate activations represented by $\vec{\mathbf{a}}^\ell$. The latter performs the same computation in the bottom-up direction and are represented by $\overleftarrow{\mathbf{a}}^\ell$. The

aggregation parameters \overrightarrow{U}^ℓ and \overleftarrow{U}^ℓ are defined for combining both directions and obtaining final activations \mathbf{a}^ℓ for concept layer ℓ .

The entire formulation for a feature vector \mathbf{x} , after obtaining \mathbf{x}^ℓ for all concept layers using Eq. (3.1), is the following:

$$\overrightarrow{\mathbf{a}}^\ell = \overrightarrow{V}^{\ell-1,\ell} \cdot \overrightarrow{\mathbf{a}}^{\ell-1} + \overrightarrow{H}^\ell \cdot \mathbf{x}^\ell + \overrightarrow{\mathbf{b}}^\ell, \quad (3.2)$$

$$\overleftarrow{\mathbf{a}}^\ell = \overleftarrow{V}^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}^{\ell+1} + \overleftarrow{H}^\ell \cdot \mathbf{x}^\ell + \overleftarrow{\mathbf{b}}^\ell, \quad (3.3)$$

$$\mathbf{a}^\ell = \overrightarrow{U}^\ell \cdot \overrightarrow{\mathbf{a}}^\ell + \overleftarrow{U}^\ell \cdot \overleftarrow{\mathbf{a}}^\ell + \mathbf{b}^{a,\ell}, \quad (3.4)$$

where $V^{i,j} \in \mathbb{R}^{n_j \times n_i}$, $H^i \in \mathbb{R}^{n_i \times n_i}$, $U^i \in \mathbb{R}^{n_i}$, and $\mathbf{b}^{a,i}, \mathbf{b}^i \in \mathbb{R}^{n_i}$ are all learnable parameters. The V's and H's weights capture the inter-layer and intra-layer dependencies, respectively. Since these parameters exhaust all pairwise relationships between labels, this step can be thought of as propagating activations across a fully-connected directed label graph.

In order to obtain concept-specific probabilities for making predictions, the activations a_k^ℓ are passed through a sigmoid function (i.e. $\sigma(z) = \frac{1}{1+e^{-z}}$), for concepts $k \in \ell$ and concept layers $\ell \in \mathcal{M}$, yielding normalized activations $y_k^\ell = \sigma(a_k^\ell)$, used as probabilities. To learn the layer parameters, the model is trained end-to-end with backpropagation, minimizing logistic cross-entropy loss for a given batch of activations $\{a_k^\ell\}$, as follows:

$$E(\{a_k^\ell\}) = - \sum_{i=1}^N \sum_{\ell=1}^m \sum_{k=1}^{n_\ell} \left(t_{k,i}^\ell \cdot \log(\sigma(a_{k,i}^\ell)) + (1 - t_{k,i}^\ell) \cdot \log(1 - \sigma(a_{k,i}^\ell)) \right), \quad (3.5)$$

where N denotes the batch size selected as a hyperparameter, $t_{k,i}^\ell$ and $\sigma(a_{k,i}^\ell)$ correspond to the ground-truth label and output score, respectively, for the k -th concept in layer ℓ for the i -th sample in the batch.

3.2 Structured Inference Neural Network (SINN)

The fully-connected bidirectional model is capable of representing all types of label relations. In practice, however, it is hard to train a model on limited data due to the large number of free parameters. To avoid this problem, we use a structured label relation graph to restrict information propagation.

We use structured label relations of positive correlation and negative correlation as prior knowledge to refine the model. The intuition is as follows: since we know that *office* is an *indoor* scene, *beach* is an *outdoor* scene, and *indoor* and *outdoor* are mutually exclusive, a high score on *indoor* should increase the probability of label *office* and decrease the probability of label *beach*. Labels that are not semantically related, e.g. *motorcycle* and *shoebox*, should not affect each other. The structured label relations can be obtained from

sources such as semantic taxonomies, or by parsing WordNet relations [36]. We introduce the notation V_p , V_n , H_p and H_n to explicitly capture structured label relations in between and within concept layers, where the subscripts p and n indicate positive and negative correlation, respectively. These model parameters are masked matrices capturing the label relations. Instead of learning full parametrized matrices V_p , V_n , H_p and H_n , we freeze some elements to be zero if there is no semantic relation between the corresponding labels. For example, V_p models the positive correlation in between two concept layers: only the label pairs that have positive correlation have learnable model parameters, while the rest are zeroed out to remove potential noise. A similar setting goes to V_n , H_p and H_n . Figure 3.2 shows an example positive correlation graph and a negative graph between two layers.

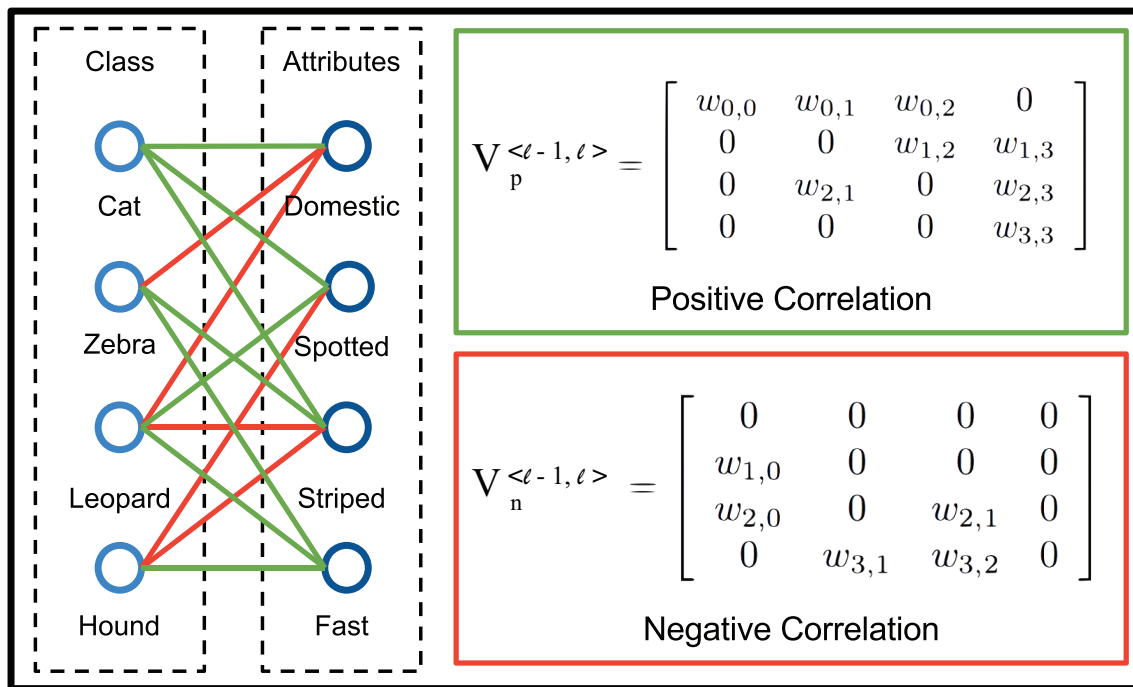


Figure 3.2: An example showing the model parameters V_p and V_n between the animal layer and the attribute layer. Green edges in the graph represent positive correlation, and red edges represent negative correlation.

To implement the positive and negative label correlation, we implement the following structured message passing process:

$$\begin{aligned} \vec{\mathbf{a}}^\ell &= \gamma(\vec{V}_p^{\ell-1,\ell} \cdot \vec{\mathbf{a}}^{\ell-1}) + \gamma(\vec{H}_p^\ell \cdot \mathbf{x}^\ell) \\ &\quad - \gamma(\vec{V}_n^{\ell-1,\ell} \cdot \vec{\mathbf{a}}^{\ell-1}) - \gamma(\vec{H}_n^\ell \cdot \mathbf{x}^\ell) + \vec{\mathbf{b}}^\ell, \end{aligned} \quad (3.6)$$

$$\begin{aligned} \overleftarrow{\mathbf{a}}^\ell &= \gamma(\overleftarrow{V}_p^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}^{\ell+1}) + \gamma(\overleftarrow{H}_p^\ell \cdot \mathbf{x}^\ell) \\ &\quad - \gamma(\overleftarrow{V}_n^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}^{\ell+1}) - \gamma(\overleftarrow{H}_n^\ell \cdot \mathbf{x}^\ell) + \overleftarrow{\mathbf{b}}^\ell, \end{aligned} \quad (3.7)$$

$$\mathbf{a}^\ell = \vec{U}^\ell \cdot \vec{\mathbf{a}}^\ell + \overleftarrow{U}^\ell \cdot \overleftarrow{\mathbf{a}}^\ell + \mathbf{b}^{a,\ell}. \quad (3.8)$$

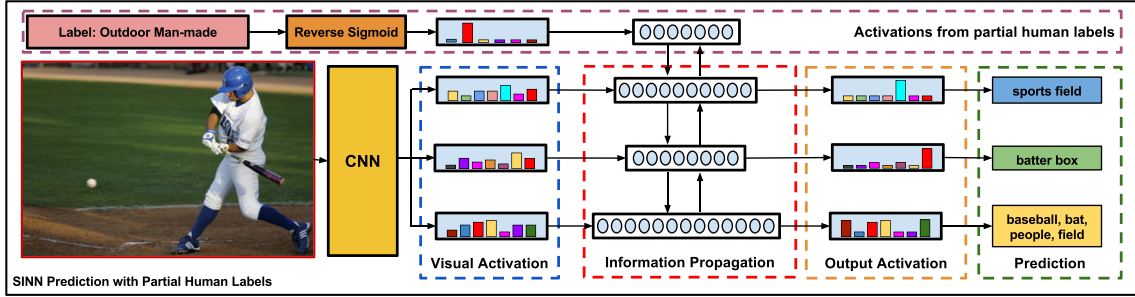


Figure 3.3: The label prediction pipeline with partial observation. The pipeline is similar to Figure 3.1 except that we now have a partial observation that this image is outdoor man-made. The SINN is able to take the observed label into consideration and improve the label predictions in the other concept layers.

Here $\gamma(\cdot)$ stands for a ReLU activation function. It is essential for SINN as it enforces that activations from positive correlation always make positive contribution to output activation and keeps activations from negative correlation as negative contribution (notice the minus signs in Eqs. (3.6) and (3.7)). To learn the model parameters V 's, H 's, and U 's, we optimize the cross-entropy loss in Eq. (3.5).

3.3 Modeling Cross-temporal Relations

In order to model temporal dependencies between frames benefiting from the inclusion of structured label relationships, we extend BINN to allow message passing across time. Since this temporal extension is built on top of a LSTM model (Chapter 2.6) and benefits from bidirectional inference in label space, we refer to it as bidirectional inference LSTM (biLSTM) from now on. Thus, considering an input sequence $\{\mathbf{x}_t\}_{t=1}^T$, we can rewrite Eq. (3.1) as follows:

$$\mathbf{x}_t^\ell = W^\ell \cdot \mathbf{x}_t + \mathbf{b}^\ell, \quad (3.9)$$

where W^ℓ and \mathbf{b}^ℓ are parameters shared across time.

Using the same notation for the bidirectional message passing presented in Chapter 3.1, Eq. (3.2), (3.3) and (3.4) can be extended as follows:

$$\vec{\mathbf{a}}_t^\ell = \vec{V}^{\ell-1,\ell} \cdot \vec{\mathbf{a}}_t^{\ell-1} + \vec{H}^\ell \cdot \mathbf{x}_t^\ell + \vec{\mathbf{b}}^\ell, \quad (3.10)$$

$$\overleftarrow{\mathbf{a}}_t^\ell = \overleftarrow{V}^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}_t^{\ell+1} + \overleftarrow{H}^\ell \cdot \mathbf{x}_t^\ell + \overleftarrow{\mathbf{b}}^\ell, \quad (3.11)$$

$$\mathbf{a}_t^\ell = \vec{U}^\ell \cdot \vec{\mathbf{a}}_t^\ell + \overleftarrow{U}^\ell \cdot \overleftarrow{\mathbf{a}}_t^\ell + \mathbf{b}^{a,\ell}, \quad (3.12)$$

where $V^{i,j} \in \mathbb{R}^{n_j \times n_i}$, $H^i \in \mathbb{R}^{n_i \times n_i}$, $U^i \in \mathbb{R}^{n_i}$, and $\mathbf{b}^{a,i}, \mathbf{b}^i \in \mathbb{R}^{n_i}$ are shared across time.

Next, we combine the output \mathbf{a}_t^ℓ with the hidden state of a LSTM after each time step. Assuming $\sigma(\cdot)$ and $\tanh(\cdot)$ are the sigmoid and hyperbolic tangent functions, we rewrite the LSTM unit definition from [57], as follows:

$$\mathbf{i}_t^\ell = \sigma(W_{xi}^\ell \cdot \mathbf{x}_t + W_{hi}^\ell \cdot \mathbf{h}_{t-1}^\ell + \mathbf{b}_i^\ell), \quad (3.13)$$

$$\mathbf{f}_t^\ell = \sigma(W_{xf}^\ell \cdot \mathbf{x}_t + W_{hf}^\ell \cdot \mathbf{h}_{t-1}^\ell + \mathbf{b}_f^\ell), \quad (3.14)$$

$$\mathbf{o}_t^\ell = \sigma(W_{xo}^\ell \cdot \mathbf{x}_t + W_{ho}^\ell \cdot \mathbf{h}_{t-1}^\ell + \mathbf{b}_o^\ell), \quad (3.15)$$

$$\mathbf{g}_t^\ell = \tanh(W_{xc}^\ell \cdot \mathbf{x}_t + W_{hc}^\ell \cdot \mathbf{h}_{t-1}^\ell + \mathbf{b}_c^\ell), \quad (3.16)$$

$$\mathbf{c}_t^\ell = \mathbf{f}_t \odot \mathbf{c}_{t-1}^\ell + \mathbf{i}_t^\ell \odot \mathbf{g}_t^\ell, \quad (3.17)$$

$$\mathbf{h}_t^\ell = \mathbf{o}_t^\ell \odot \tanh(\mathbf{c}_t^\ell), \quad (3.18)$$

where \odot denotes point-wise multiplication. This formulation is equivalent to appending one LSTM unit for each concept layer ℓ . Intuitively, we allow the model to accumulate information about each concept layer separately, which empirically provided better results than assigning one memory unit for all concept layers.

The concept layer's activations \mathbf{a}_t^ℓ are then combined with its memory unit's hidden state \mathbf{h}_t^ℓ at time step t , as follows:

$$\mathbf{y}_t^\ell = \sigma(M^a \cdot \mathbf{a}_t^\ell + M^h \cdot \mathbf{h}_t^\ell + \mathbf{b}^{a,h}), \quad (3.19)$$

where M^a , M^h and $\mathbf{b}^{a,h}$ are parameters shared across time, \mathbf{y}_t^ℓ corresponds to normalized activations for concept layer ℓ at time step t used as confidence scores and $\sigma(\cdot)$ is the sigmoid function.

Similarly, we extend the SINN model to allow the flow of information across time as well, by rewriting Eqs. (3.6) and (3.7), as follows:

$$\begin{aligned} \vec{\mathbf{a}}_t^\ell &= \gamma(\vec{V}_p^{\ell-1,\ell} \cdot \vec{\mathbf{a}}_t^{\ell-1}) + \gamma(\vec{H}_p^\ell \cdot \mathbf{x}_t^\ell) \\ &\quad - \gamma(\vec{V}_n^{\ell-1,\ell} \cdot \vec{\mathbf{a}}_t^{\ell-1}) - \gamma(\vec{H}_n^\ell \cdot \mathbf{x}_t^\ell) + \vec{\mathbf{b}}^\ell, \end{aligned} \quad (3.20)$$

$$\begin{aligned} \overleftarrow{\mathbf{a}}_t^\ell &= \gamma(\overleftarrow{V}_p^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}_t^{\ell+1}) + \gamma(\overleftarrow{H}_p^\ell \cdot \mathbf{x}_t^\ell) \\ &\quad - \gamma(\overleftarrow{V}_n^{\ell+1,\ell} \cdot \overleftarrow{\mathbf{a}}_t^{\ell+1}) - \gamma(\overleftarrow{H}_n^\ell \cdot \mathbf{x}_t^\ell) + \overleftarrow{\mathbf{b}}^\ell, \end{aligned} \quad (3.21)$$

using the same notation as for the previously formulated models. Note that the aggregation equation stays the same as in Eq. (3.12). In order to distinguish this model from the biLSTM, we refer to it as structured inference LSTM (siLSTM). The output activations \mathbf{y}_t^ℓ serve as confidence scores for each concept at layer ℓ being assigned to the frame at time step t and is used for outputting the detections.

3.4 Prediction Framework

Now we introduce the method of predicting labels during test time. As the model is trained with multiple concept layers, it is straightforward to recognize a label at each concept layer for the provided test sample. This mechanism is called label prediction *without observation* (the default pipeline shown in Figure 3.1).

Another interesting application is to make predictions with *partial observations* – we want to predict labels in one concept layer given labels in another concept layer. Figure 3.3 illustrates the idea. Given an image shown in the left side of Figure 3.3, we have more confidence to predict it as *batter’s box* once we know it is an *outdoor* image with attribute *sports field*.

To make use of the partially observed labels in our SINN framework, we need to transform the observed binary labels into soft activation scores for SINN to improve the label prediction on the target concept layers. Recall that SINN minimizes cross-entropy loss which applies sigmoid functions on activations to generate label confidences. Thus, we reverse this process by applying the inverse sigmoid function (logit) on the binary ground-truth labels to obtain activations. Formally, we define the k -th observed concept activation a_k^ℓ in concept layer ℓ obtained from a ground-truth label t_k^ℓ as:

$$a_k^\ell = \begin{cases} \log \frac{t_k^\ell + \epsilon}{1 - (t_k^\ell + \epsilon)}, & \text{if } t_k^\ell = 0, \\ \log \frac{t_k^\ell - \epsilon}{1 - (t_k^\ell - \epsilon)}, & \text{if } t_k^\ell = 1. \end{cases} \quad (3.22)$$

Note that we put a small perturbation ϵ on the ground-truth label t_k^ℓ for numerical stability. In our experiments, we set $\epsilon = 0.001$.

Chapter 4

Multi-label Classification on Images

We start by describing our evaluation benchmarks for the task of multi-label classification on images (Chapter 4.1) as well as the corresponding metrics used for evaluation (Chapter 4.2). Next, we present our experimental setup and implementation details in Chapter 4.3 and 4.4, respectively. Finally, we report the results for this task in Chapter 4.5.

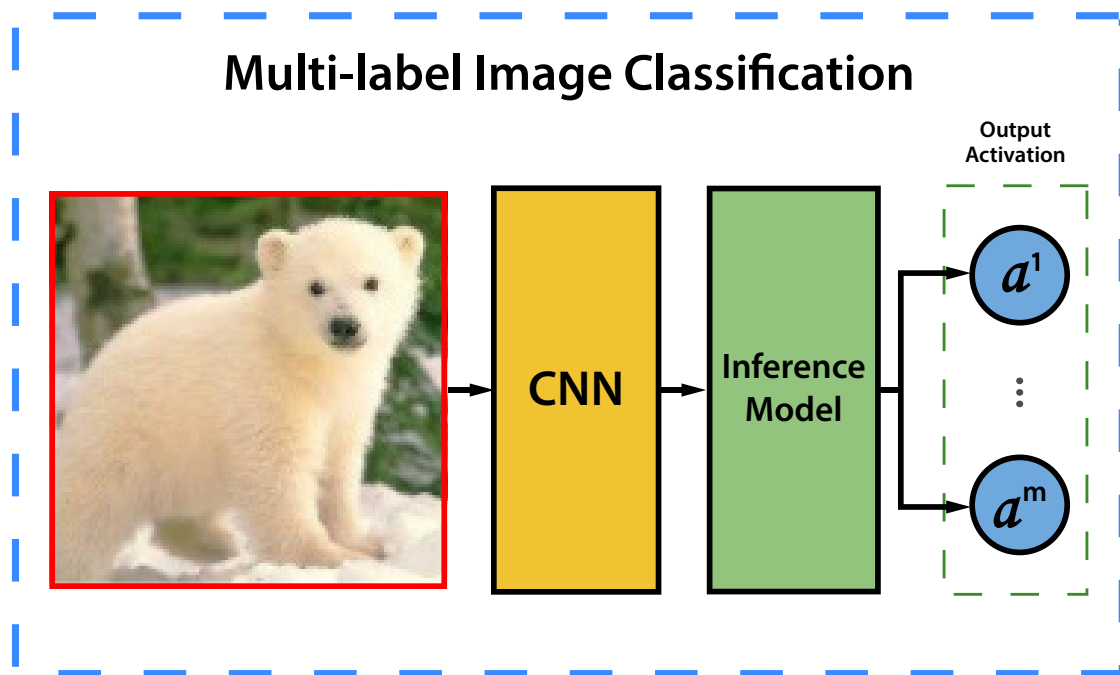


Figure 4.1: The inference model is applied directly on a feature vector extracted by a CNN for performing multi-label image classification.

4.1 Evaluation Benchmarks

4.1.1 Animal with Attributes (AwA)

The *Animals with Attributes* (AwA) dataset was first introduced in [30] for validating attribute-based classification performed in a zero-shot fashion. The AwA dataset consists of 30,475 images of animals queried from image search engines of *Microsoft*, *Google*, *Yahoo* and *Flickr*. The annotations comprises 50 animal classes and 85 and semantic attributes. In addition to the original annotations a 28-taxonomy-term layer was included [15] by parsing WordNet.

4.1.2 Scene Understanding (SUN)

The *Scene Understanding* (SUN) database [54] contains a total of 908 categories and 131,072 images. The final set of scene categories was obtained by parsing WordNet for collecting commonly found scenes. For the task of scene categorization only 397 categories (SUN397), for which there are at least 100 unique photographs were used for benchmarking. The SUN397 database contains an intrinsic hierarchical structure in label space defined by three levels of granularity annotated by Amazon’s Mechanical Turk (AMT) workers. In increasing order of granularity SUN397 (from top to bottom) contains on the top layer 3 coarse categories, 16 general scene categories and 397 fine-grained scene categories on the bottom layer.

4.1.3 Real-World Web Images (NUS-WIDE)

The original NUS-WIDE [3] dataset consists of a large scale web image collection containing unique associated tags from Flickr checked against WordNet concepts. Over 300,000 images were randomly crawled together with their tags through the Flickr public API. After filtering duplicates, a total of 269,648 images remain, which include 5,018 associated unique tags and 81 ground-truth concepts.

4.2 Evaluation Metrics

Here we define the metrics used for the multi-label image classification. All formulations are derived for a collection of images \mathcal{I} with candidate classes \mathcal{C} . For each image sample $I \in \mathcal{I}$ with a ground-truth vector \mathbf{t}_i are assigned multiple labels represented by a vector \mathbf{y}_i .

4.2.1 Multiclass Accuracy (MC Acc)

The *MC Acc* metric is a standard metric used for image classification and consists of averaging the per-class accuracies for all classes in the problem:

$$MC\ Acc = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \frac{TP_c(\mathcal{I})}{|\mathcal{I}|}, \quad (4.1)$$

where $TP_c(\mathcal{I})$ denotes the true positives predictions for class c , considering all images in \mathcal{I} .

4.2.2 Intersection over Union Accuracy (IoU Acc)

Another widely used metric is the *IoU Acc*, which is based on the *Hamming distance* between predictions and the ground-truth labels:

$$IoU\ Acc = \frac{1}{|\mathcal{I}|} \sum_{i=1}^{|\mathcal{I}|} HammingDistance(\mathbf{y}_i, \mathbf{t}_i). \quad (4.2)$$

4.2.3 Average Precision (AP)

In order to summarize the area under the precision-recall curve, the AP can be computed by averaging the precision weighed by increments in recall at each element in a list of annotation scores ranked by the model scores, as follows:

$$AP = \sum_{i=1}^N P(i) \Delta R(i), \quad (4.3)$$

where N is the number of predictions, $P(i)$ and $\Delta R(i)$ are precision and increment recall at the i -th element in the ranked list.

4.2.4 Per-label Mean Average Precision (mAP_L)

The mAP_L is obtained by computing separately, the average precision (AP) for each label across all the samples in the target dataset. Given a list of annotation scores ranked by the model scores, we compute AP_c for each class using Eq. (4.3) and calculate the mean over all classes:

$$mAP_L = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} AP_c. \quad (4.4)$$

4.2.5 Per Image Mean Average Precision (mAP_I)

As opposed to mAP_L, the mAP_I considers the ranking of predictions for each image sample and compute the average across all samples in the dataset:

$$mAP_I = \frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} AP_I, \quad (4.5)$$

where AP_I is computed using Eq. (4.3) for each image disregarding the classes.

4.3 Experimental Setup

The default validation framework on the image datasets (i.e. AwA, NUS-WIDE and SUN397) consists of comparing our method against baselines and state-of-the-art methods. In addition, for all experiments the spatial feature extraction is done by applying a pre-trained CNN underneath the classification algorithm.

4.3.1 AwA: Layered Prediction with Label Relations

This experiment demonstrates the label prediction capability of our SINN model and the effectiveness of adding structured label relations for label prediction. We run each method five times with five random splits – 60% for training and 40% for test. We report the average performance as well as the standard deviation of each performance measure. Note that there is very little related work with layered label prediction on AwA. The most relevant one is work by Hwang and Sigal [17] on unified semantic embedding (USE). The comparison is not strictly fair, as the train/test splits are different. Further, we include our BINN model without specifying the label relation graphs (see Chapter 3.1) as a baseline method in this experiment, as it can verify the performance gain in our model from including structure.

4.3.2 NUS-WIDE: Multi-label Classification with Partial Human Labels of Tags and Groups

This experiment shows our model’s capability to use noisy tags and structured tag-label relations to improve multi-label classification. As previous work used various evaluation metrics and experiment settings, and there are no fixed train/test splits, it is hard to make direct comparisons. Also note that a fraction of previously used images are unavailable now due to Flickr copyright.

In order to make our result as comparable as possible, we tried to set up the experiments according to previous work. We collected all available images and discard images with missing labels as previous work did [9, 21], and got 168,240 images of the original dataset. To make our result comparable with [21], we use 5 random splits with the same train/test ratio as [21] – there are 132,575 training images and 35,665 test images in each split. To compare our method with [21, 35], we also used the tags and metadata groups in our experiment. Different from their settings, instead of augmenting images with 5000 tags, we only used 1000 tags, and augment the image with 698 group labels obtained from image metadata to form a three-layer group-concept-tag graph. Instead of using the tags as sparse binary input features (as in [21, 35]), we convert them to observed labels and feed them to our model.

We report our results on this dataset with two settings for our SINN, the first using 1k tags as the only observations to a bottom level of the relation graph. This method provides a good comparison to the tag neighborhood + tag vector [14], as we did not use

extra information other than tags. In the second setting, we make both group and tag levels observable to our SINN, which achieves the best performance. We also compared our results with that of McAuley *et al.* [35], Gong *et al.* [9].

4.3.3 SUN397: Improving Scene Recognition with and without partially Observed Labels

We conducted two experiments on the SUN397 dataset. The first experiment is similar to the study on AwA: we applied our model to layered image classification with label relations, and compare our model with CNN + Logistics and CNN + BINN baselines, as well as a state-of-the-art approach [55, 54]. For fair comparison, we used the same train/test split ratio as [55, 54], where we have 50 training and test images in each of the 397 scene categories. To migrate the randomness in sampling, we also repeat the experiment 5 times and report the average performance as well as the standard deviations. In the second experiment, we considered partially observed labels from the top (coarsest) scene layer as input to our inference framework. In other words, we assume we know whether an image is *indoor*, *outdoor man-made*, or *outdoor natural*.

4.4 Implementation Details

To optimize our learning objective, we use stochastic gradient descent with mini-batch size of 50 images and momentum of 0.9. For all training runs, we apply a two-stage policy as follows. In the first stage, we fixed pretrained CNN networks, and train our SINN with a learning rate of 0.01 with fixed-size decay step. In the second stage, we set the learning rate as 0.0001 and fine-tune the CNN together with our SINN. We set the gradient clipping threshold to be 25 to prevent gradient explosion. The weight decay value for our training procedure is set to 0.0005. In the computation of visual activations from the CNN, as different experiment datasets describe different semantic domains, we adopt different pretrained CNN models: ImageNet pretrained model [19] for experiments 4.3.1 and 4.3.2, placenet pretrained model [58] for experiment 4.3.3.

4.5 Experimental Results

For all image benchmarks, the experimental results show that the inclusion of label relation graphs effectively boost performance on the three datasets tested in [15] (i.e. AwA, SUN397 and NUS-WIDE). From Table 4.1, 4.2 and 4.4, we can see that SINN consistently improves performance significantly in most experiments executed.

Table 4.1 shows that our method outperforms the baseline methods (CNN + Logistics and CNN + BINN variants) as well as the USE method, in terms of each concept layer and each performance metric. It validates the efficacy of our proposed model for image

Concept Layer	Model	MC Acc	IoU	mAP_L
28 taxonomy terms	CNN + Logistics	-	80.41±0.09	90.16±0.10
	CNN + BINN	-	79.85±0.13	89.92±0.07
	CNN + SINN	-	84.47±0.38	93.00±0.29
50 animal classes	USE [17] + DECAF [8]	46.42 ±1.33	-	-
	CNN + Logistics	78.44±0.27	62.75±0.26	78.35±0.19
	CNN + BINN	79.00±0.43	62.80±0.25	78.88±0.35
	CNN + SINN	79.36±0.43	66.60±0.43	81.19±0.14
85 attributes	CNN + Logistics	-	81.29±0.10	93.29±0.12
	CNN + BINN	-	80.64±0.13	93.04±0.13
	CNN + SINN	-	86.92±0.18	96.05±0.07

Table 4.1: Layered label prediction results on the AwA dataset.

classification. Note that for the results in Table 1, we did not finetune the first seven layers of the CNN [28] for fairer comparison with Hwang and Sigal [16] (which only makes use of DECAF features [8]). Fine-tuning the first seven CNN layers further improves IoU_{Acc} at each concept layer to 86.06 ± 0.72 (28 taxonomy terms), 69.17 ± 1.00 (50 animal classes), 88.22 ± 0.38 (85 attributes), and mAP_L to 94.17 ± 0.55 (28 taxonomy terms), 83.12 ± 0.69 (50 animal classes), 96.72 ± 0.20 (85 attributes), respectively.

The results on SUN397 are summarized in Table 4.2, showing that our proposed method again achieves a considerable performance gain over all the compared methods. In Table 4.3, we compare the 397 fine-grained scene recognition performance. We compare to a set of baselines, including CNN + Logistics + Partial Labels that considers the partial labels as an extra binary indicator feature vector for logistic regression. Results show that our method combined with partial labels (i.e., CNN + SINN + Partial Labels) improves over baselines, exceeding the second best by 4% MC_{Acc} and 6% mAP_L .

The results on NUS-WIDE are shown in Table 4.4. We can see that SINN outperformed all the baseline methods and existing approaches (i.e.[21, 9, 35]) by a large margin on NUS-WIDE dataset, considering different settings. The results on SUN397 are summarized in Table 4.4 and present consistent improvements when using SINN.

Concept Layer	Model	MC Acc	IoU	mAP_L
3 coarse scene categories	CNN + Logistics	-	83.67±0.18	95.19±0.07
	CNN + BINN	-	83.63±0.24	95.19±0.03
	CNN + SINN	-	85.95±0.44	96.40±0.18
16 general scene categories	CNN + Logistics	-	64.30±0.27	83.30±0.19
	CNN + BINN	-	63.40±0.35	82.93±0.14
	CNN + SINN	-	66.46±1.10	84.97±0.96
397 fine-grained scene categories	Image features + SVM [55, 54]	42.70	-	-
	CNN + Logistics	57.86±0.18	35.97±0.37	55.31±0.30
	CNN + BINN	57.52±0.29	35.44±1.02	55.57±0.63
	CNN + SINN	57.60±0.38	37.71±1.13	58.00±0.33

Table 4.2: Layered label prediction results on the SUN397 dataset.

Model	MC+Acc	mAP_L
Image features + SVM [55, 54]	42.70	-
CNN + Logistics	57.86±0.38	55.31±0.30
CNN + BINN	57.52±0.29	55.57±0.63
CNN + SINN	57.60±0.38	58.00±0.33
CNN + Logistics + Partial Labels	59.08±0.27	56.88±0.29
CNN + SINN + Partial Labels	63.46±0.18	64.63±0.28

Table 4.3: Recognition results on the 397 fine-grained scene categories. Note that the last two compared methods make use of partially observed labels from the top (coarsest) scene layer, i.e. *indoor*, *outdoor man-made*, and *outdoor natural*.

Model	mAP _L	mAP _I	R _L	P _L	R _I	P _I
Graphical Model [35]	49.00	-	-	-	-	-
CNN + WARP [9]	-	-	35.60	31.65	60.49	48.59
5k tags + Logistics [21]	43.88 ±0.32	77.06 ±0.14	47.52 ±2.59	46.83 ±0.89	71.34 ±0.16	51.18 ±0.16
5k tags + Tag neighbors [21]	61.88 ±0.36	80.27 ±0.08	57.30 ±0.44	54.74 ±0.63	75.10 ±0.20	53.46 ±0.09
CNN + Logistics	46.94 ±0.47	72.25 ±0.19	45.03 ±0.44	45.60 ±0.35	70.77 ±0.21	51.32 ±0.14
1k tags + Logistics	50.33 ±0.37	66.57 ±0.12	23.97 ±0.23	47.40 ±0.07	64.95 ±0.18	47.40 ±0.07
1k tags + Groups + Logistics	52.81 ±0.40	68.04 ±0.12	25.54 ±0.24	49.26 ±0.15	65.99 ±0.15	48.13 ±0.05
1k tags + Groups + CNN + Logistics	54.67 ±0.57	77.81 ±0.22	50.83 ±0.53	49.36 ±0.30	75.38 ±0.16	54.61 ±0.09
1k tags + CNN + SINN	67.20 ±0.60	81.99 ±0.14	59.82 ±0.12	57.02 ±0.57	78.78 ±0.13	56.84 ±0.07
1k tags + Groups + CNN + SINN	69.24 ±0.47	82.53 ±0.15	60.63 ±0.67	58.30 ±0.33	79.12 ±0.18	57.05 ±0.09

Table 4.4: Results on NUS-WIDE for from [15]. We also included the results for per-label recall and precision (R_L and P_L) and per-image (R_I and P_I), considering the top 3 labels for each image.

Chapter 5

Multi-label Classification on Videos

In this section, we describe the work done for multi-label classification on videos. In Chapter 5.1 and 5.2, we present details on the two YouTube-8M benchmarks as well as the metrics used for this task, respectively. Chapter 5.3 and 5.4, present our experimental setup and implementation details, in this order. Lastly, we report the results obtained in Chapter 5.5.

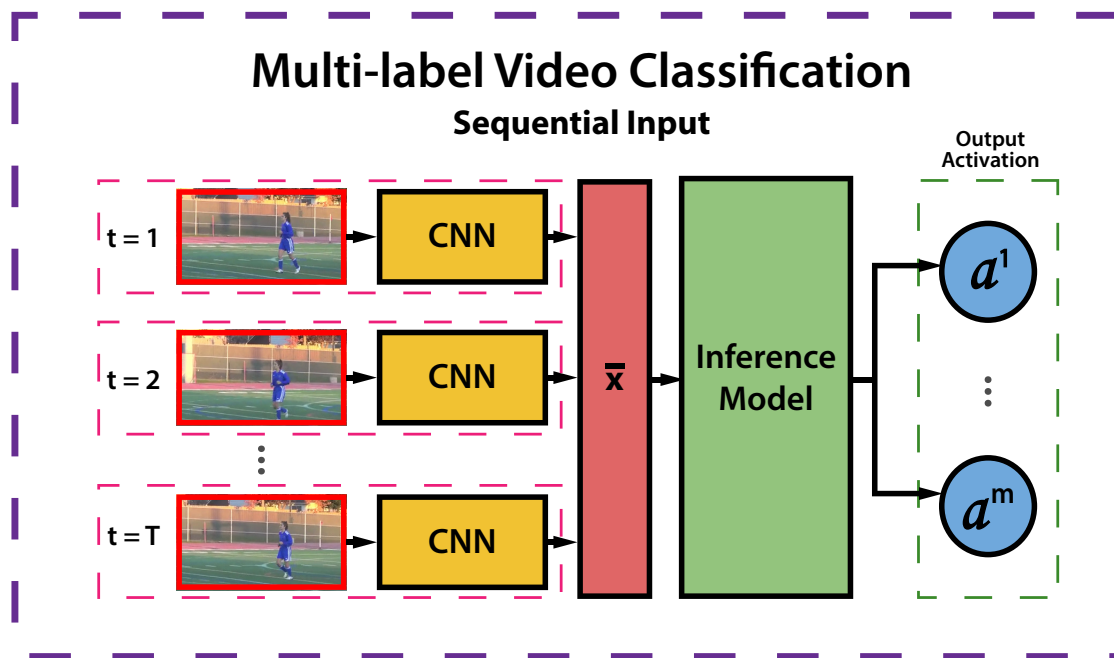


Figure 5.1: The video-level representation $\bar{\mathbf{x}}$ is obtained by pooling per-frame feature vectors \mathbf{x}_t (using a CNN) and fed to the inference model for performing multi-label video classification.

5.1 Evaluation Benchmarks

5.1.1 YouTube-8M (YT-8M)

The YouTube-8M dataset consists of approximately 8 million YouTube videos, each annotated with 4800 Google Knowledge Graph *entities*, functioning as classes. With each entity label is associated up to 3 *verticals* (i.e. coarse-grained labels). The dataset is derived from roughly 500K hours of untrimmed videos, with an average of 1.8 labels per video. Each video is decoded as a set of features extracted by passing the RGB frame through the InceptionV3 model from Szegedy *et al.* [48], a deep CNN pretrained on ImageNet [41], and Principal Component Analysis (PCA) is applied to reduce feature dimension. The scale of this dataset in both label space and data space is unprecedented in the field of video datasets, surpassing previous benchmarks such as UCF-101 and Sports1M.

5.1.2 YouTube-8M V2 (YT-8M V2)

The YouTube-8M V2 dataset represents the frame and audio features from approximately 7 million YouTube videos. The dataset is an updated version of YouTube-8M, with an increased number of labels per video and a smaller number of entities. On average, the videos in YT-8M V2 have 3.4 labels each, and there are only 4716 Google Knowledge Graph entities forming the label space. The preprocessing for this dataset is the same as YT-8M, but the audio features are also included, calculated using the CNN method in [13].

5.2 Evaluation Metrics

The metrics formulated in this section considers a collection of videos \mathcal{V} , such that for each video $v \in \mathcal{V}$, a model predicts a set of labels $\mathcal{P}_v \in \mathcal{C}$ and the metric is computed against a set of ground-truth labels $\mathcal{G}_v \in \mathcal{C}$. We start by defining the indicator function for a generic set \mathcal{A} and an element x , used in further formulations, as follows:

$$\mathbb{1}_{\mathcal{A}}(x) = \begin{cases} 1, & \text{if } x \in \mathcal{A}, \\ 0, & \text{otherwise.} \end{cases} \quad (5.1)$$

5.2.1 Hit at Top K (Hit@k)

The *Hit@k* corresponds to the portion of samples whose top k predictions contain at least one ground-truth label and can be written as follows:

$$Hit@k(\mathcal{V}) = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \vee_{e \in top(\mathcal{P}_v, k)} \mathbb{1}_{\mathcal{G}_v}(e), \quad (5.2)$$

where \vee is the logical OR operator.

5.2.2 Precision at Equal Recall Rate (PERR)

The definition for this metric is the same as in [1] and consists of the precision computed for each sample when retrieving the same number of labels as the ground-truth labels:

$$PERR(\mathcal{V}) = \frac{1}{|\mathcal{V}'|} \sum_{v \in \mathcal{V}'} \left[\frac{1}{|\mathcal{G}_v|} \sum_{e \in \text{top}(\mathcal{P}_v, |\mathcal{G}_v|)} \mathbf{1}_{\mathcal{G}_v}(e) \right], \quad (5.3)$$

where $\mathcal{V}' \subseteq \mathcal{V}$ is the subset of test videos containing at least one ground-truth label.

5.2.3 Global Average Precision (gAP)

The gAP metric defined in [1] is computed similarly to the mAP_L (Chapter 4.2.4) defined earlier, with the distinction that the gAP is computed agnostically to the top k (i.e. k=20) classes as for a binary problem:

$$gAP = \sum_{i=1}^{kN_v} P(i) \Delta R(i), \quad (5.4)$$

where N_v is the number of videos in the test set, $P(i)$ and $\Delta R(i)$ are defined as for Eq. (4.3)

5.2.4 Per Video Mean Average Precision (mAP_V)

As is [1], buckets of length 10^{-4} are used for discretizing the precision-recall threshold τ for each class. All the non-zero annotations are then sorted (ascending order) according to the prediction scores. The precision $P_c(\tau)$ and recall $R_c(\tau)$ at a given threshold τ for class c is defined as follows:

$$P_c(\tau) = \frac{\sum_{i=1}^{N_v} (y_{ic} \geq \tau) t_{ic}}{\sum_{i=1}^{N_v} (y_{ic} \geq \tau)}, \quad (5.5)$$

$$R_c(\tau) = \frac{\sum_{i=1}^{N_v} (y_{ic} \geq \tau) t_{ic}}{\sum_{i=1}^{N_v} y_{ic}}, \quad (5.6)$$

where y_{ic} and t_{ic} are prediction and ground-truth for the i -th video for class c and N_v is the total number of videos.

The AP_c for a specific class c is given by:

$$AP_c = \sum_{j=1}^{10^4} P_c(\tau_j) [R_c(\tau_j) - R_c(\tau_{j+1})], \quad (5.7)$$

where $\tau_j = \frac{j}{10^4}$. The mAP_V is obtained by taking the unweighted mean across all classes in \mathcal{C} (as in Eq. (4.4)).

5.3 Experimental Setup

In order to validate our method on multi-label video classification, we follow a similar framework to the one used for multi-label image classification (Chapter 4.3). Firstly, we reproduce the logistic baseline from [1] and compare against our models for both releases of YouTube-8M. For YT-8M V2, we conduct experiments with spatial features only and spatial plus audio features, comparing our method against the logistic baseline. Additionally, following the same intuition from the partially observed experiments for multi-label image classification. We run experiments providing the partial observations of the verticals for both datasets to our model and validate against the logistic baseline including the partial observations as input features. Note that including partial observations to the logistic regression is important to demonstrate that partial labels are not the discriminative feature for these experiments. All experiments were executed following the guidelines presented in [1] and the official Kaggle competition.

5.4 Implementation Details

The dataset labels were organized into a fully-connected graph (intra- and inter-layer connections) with two concept layers – entities (i.e. fine-grained labels), and verticals (i.e. coarse-grained labels). We minimize the cross-entropy loss function using the Adam optimizer [25]. For all models, mini-batches of size 1024 were used, and a weight decay of 10^{-8} was applied. The logistic regression model was trained for 35k iterations with a learning rate of 0.01, and the BINN was trained for 90k iterations, starting with a learning rate of 0.001 with a decay factor of 0.1 at every 40k iterations. All models were implemented with the Caffe deep learning framework [19].

5.5 Experimental Results

In Table 5.1, results on YT-8M are presented on the validation set, where the first group of results refers to the baseline models reported in [1]. To ensure consistency with [1], the results for the logistic regression baseline trained by us is shown in the second group of models in Table 5.1. It is worth to mention that the inclusion of Z- and L2-normalization were crucial for duplicating the baseline results and boosting results for the other models. As can be verified from Table 5.1, the best SINN model for YT-8M achieved significant improvements on all metrics. More precisely, a 3.2% improvement was seen on mAP using exclusively video-level features, against the leading baseline. The SINN also demonstrated measurable success on the PERR, Hit@1 and gAP metrics, with respective improvements of 4.69%, 4.4% and 7.41%.

The results shown in Table 5.2 follow the same sequence as the previous results, save for the inclusion of audio features available for YT-8M V2. The best results were obtained using

Model	mAP _V	PERR	Hit@1	gAP
LSTM [1]	26.60	57.30	64.50	-
Mixture-of-Experts [1]	30.00	55.80	63.30	-
Logistics [1]	28.10	53.00	60.50	-
Logistics [38]	27.98	52.89	60.34	49.04
BINN [38]	30.17	57.40	64.48	55.76
SINN	31.18	57.58	64.74	56.39
Logistics + Partial labels	54.47	68.77	75.53	71.46
SINN + Partial labels	57.71	72.82	78.85	75.80

Table 5.1: YouTube-8M (YT-8M) results for mAP_V, Precision at Equal Recall Rate (PERR), Hit at 1 and gAP on the validation set.

Model	mAP _V	PERR	Hit@1	gAP
Logistics + RGB [38]	36.84	64.38	78.62	70.31
Logistics + RGB + Audio [38]	38.61	69.28	82.75	75.83
BINN + RGB [38]	39.12	69.01	82.27	75.99
SINN + RGB	40.19	69.11	82.35	76.33
BINN + RGB + Audio [38]	40.91	72.27	84.96	79.29
SINN + RGB + Audio	42.32	72.92	85.49	80.09
Logistics + Partial labels	50.28	78.52	89.82	84.94
SINN + Partial labels	51.88	80.55	90.99	86.51

Table 5.2: YouTube-8M V2 (YT-8M V2) results for mAP_V, Precision at Equal Recall Rate (PERR), Hit at 1 and gAP on the validation set.

RGB and audio features, with Z- and L2-Normalization. The most effective SINN model obtained 42.32%, 72.92%, 85.49% and 80.09% for mAP, PERR, Hit@1 and gAP respectively, which corresponds to improvements of 3.71%, 3.64%, 2.74% and 4.26%, against the leading baseline results. Additionally, our results indicate that our method for combining partial labels also achieves significant improvements over baselines on both releases of YouTube-8M.

Chapter 6

Dense Action Detection

Here we introduce the work done on dense action detection. Details about THUMOS'14 and MultiTHUMOS benchmarks are presented in Chapter 6.1 and metrics used for this task in Chapter 6.2. Our experimental setup is discussed in Chapter 6.3 and implementation details in Chapter 6.4. The results are reported in Chapter 6.5.

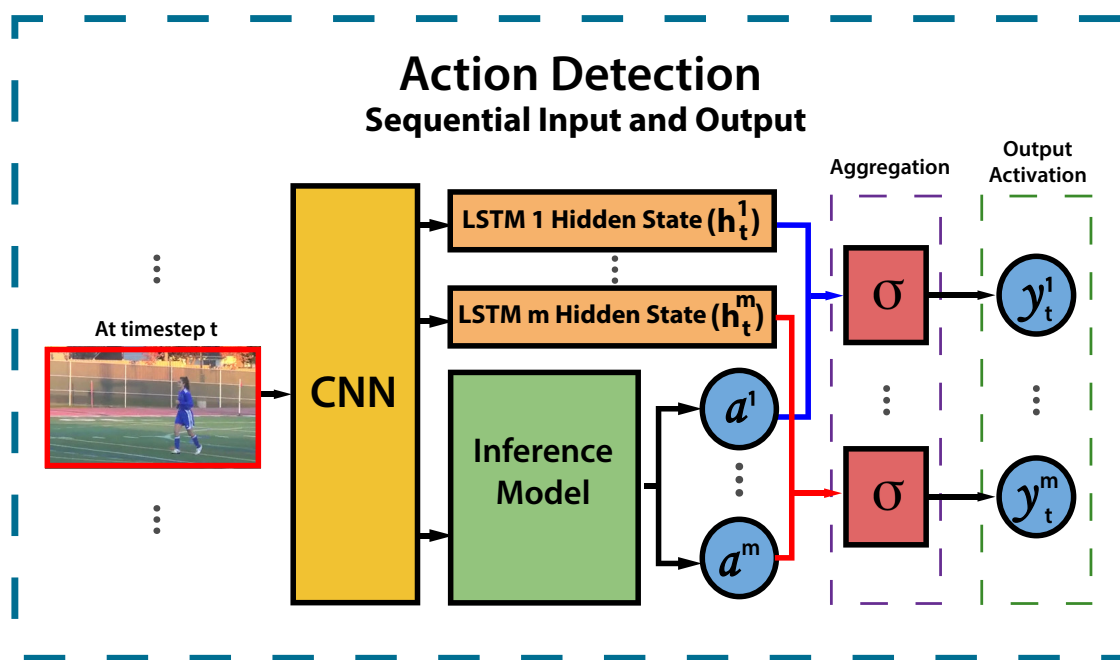


Figure 6.1: Snapshot at time step t for the action detection model, where per-frame representations \mathbf{x}_t were extracted using a CNN, fed to the inference model and to m concept-layer specific LSTM units at each time step t . The concept layer outputs $\{\mathbf{a}_t^\ell\}_{\ell=1}^m$ are combined (aggregation step denoted by σ and formulated in Eq. (3.19)) with the corresponding LSTM hidden states $\{\mathbf{h}_t^\ell\}_{\ell=1}^m$ for obtaining the final predictions $\{\mathbf{y}_t^\ell\}_{\ell=1}^m$. The output activations (y_t^ℓ) for the concept layer ℓ are interpreted as confidence scores for a given concept being assigned to the frame at time step t .

6.1 Evaluation Benchmarks

6.1.1 THUMOS'14

The THUMOS'14 [20] dataset contains 13,320 single-labeled trimmed videos from UCF101 [46] and an extra 2,584 multi-labeled untrimmed videos, from which 1,010 belong to the validation set and 1,574 to the test set. In addition to the entire set of trimmed videos, around 400 untrimmed videos (out of 2,584) are utilized for the temporal action detection task. On average these videos (i.e. trimmed and untrimmed) used for action detection, contain approximately a single label per video and 20 different action classes.

6.1.2 MultiTHUMOS

In order to study the problem of densely annotated unconstrained videos, Yeung *et al.* introduced the MultiTHUMOS dataset [57]. This dataset consists of an augmented version of the untrimmed videos for action detection from THUMOS'14, corresponding to around 30 hours of temporal action annotation. More precisely, the MultiTHUMOS dataset extends the annotation from 6,365 over 20 action classes in THUMOS'14 to 38,690 annotations over 65 action classes, including 32,325 annotations over the 45 introduced action classes. In addition, the density of annotations increases from 0.3 to 1.5 action classes per frame in average and from 1.1 to 10.5 action classes per video.

6.2 Evaluation Metrics

Analogously to Chapter 4.2.4, we use mAP_L to evaluate our method, with the difference that for this task each sample corresponds to one single frame and the mAP_L is computed across all videos as in a long global list of predictions for a binary problem. As in [57], this metric is utilized in order to avoid aggregating frame-level predictions for obtaining activities segments.

6.3 Experimental Setup

As a standard approach for this dataset, we train our model on the untrimmed validation set and validate on the test set, since the training set is the same from UCF101 and consists of short clipped videos. We executed experiments to validate the inclusion of structured annotation in two scenarios: our inference modules versus CNN in static frames and structured temporal inference versus LSTM baseline, allowing message propagation across time.

Firstly, we reproduce the single-frame CNN baseline presented in [57], to certify that we are extracting competitive features. This baseline is compared against our single-frame models, i.e. BINN and SINN. Next, we train a LSTM baseline to compare against our temporal extension models, i.e. biLSTM and siLSTM. As in [57], we report mAP_L scores

for MultiTHUMOS and the subset of classes from THUMOS on the untrimmed test set from MultiTHUMOS.

6.4 Implementation Details

The single-frame CNN baseline was pretrained on Imagenet and then fine-tuned on MultiTHUMOS. We extracted 4096-dimensional features from individual frames at fc-7 in the base model similar procedure done in [57].

In order to obtain structured annotations for our models, we augment the original set of labels \mathcal{L} by manually converting each $l \in \mathcal{L}$ to an approximate word sense from WordNet, generating a new set \mathcal{L}' . The new concept layer added to the dataset consists of selected synsets for each $l' \in \mathcal{L}'$ that surpass a certain *relatedness score* (e.g. vector score [40]) threshold (e.g. 0.5) against all synsets from WordNet. We filtered out from the new concept layer the synsets that were only related to one or most of the original labels (precise number were determined by cross-validation). We used a fully-connected graph structure with connections between all pairs of concepts for intra-layer and inter-consecutive-layer concepts.

Parsing WordNet provided us an extra layer containing 32 new concepts whose connections to the original MultiTHUMOS labels were determined by the *relatedness score*. The single-frame BINN and SINN models were trained with a batch size of 1024 frames for 8500 iterations, step size of 5000 iterations, learning rate of 10^{-3} and decay of 0.1.

For the LSTM baseline, we use the same feature vectors extracted for the single-frame CNN baseline. Differently from [57], our LSTM based models are fed with randomly sampled contiguous sequences of 32 frames (3.2s) from each video sample. In particular, the LSTM baseline was trained using ADAM with a learning rate of 10^{-3} , step size of 1500 iterations, batch size of 512 and learning rate decay of 0.1 for 1500 iterations.

Our proposed models biLSTM and siLSTM were designed to include temporal dependencies as LSTM plus, hierarchically structured labels. For these models we use the same 32 extra concepts extracted from WordNet. Also, we use a similar approach as for the LSTM model for feeding frame sequences to these models. The best model (i.e. siLSTM) obtained was trained using ADAM, batch size of 1024, learning rate of 7×10^{-4} for 900 iterations, decaying 0.1 every 700 iterations.

6.5 Experimental Results

As shown in Table 6.1, the inclusion of structured labels parsing WordNet using the single-frame SINN provides a gain of 2.7% on THUMOS and 2.4% on MultiTHUMOS against the single-frame CNN baseline implemented by us. In Figure 6.2, we show a comparison between the single-frame SINN and the single-frame CNN using per-class AP computer for each model.

Model	THUMOS	MultiTHUMOS
Single-frame CNN [57]	34.7	25.4
Single-frame CNN [Ours]	35.5	24.6
Single-frame BINN	37.6	26.8
Single-frame SINN	38.2	27.0
LSTM [57]	39.3	28.1
LSTM [Ours]	41.8	30.5
MultiLSTM [57]	41.3	29.7
biLSTM	42.8	30.6
siLSTM	43.9	31.5

Table 6.1: THUMOS and MultiTHUMOS results for mAP_L .

As observed in [57], the inclusion of temporal dependencies clearly provides substantial improvements in activity detection as we can also conclude from results shown in Table 6.1. The LSTM model implemented by us provides 6.3% THUMOS and 5.9% on MultiTHUMOS over our single-frame CNN baseline. The APs breakdown is shown in Figure 6.3.

The results from our proposed models (i.e. biLSTM and siLSTM) show some gain in performance over the LSTM baseline as show in Table 6.1. More precisely, biLSTM and siLSTM are benefiting from both inclusion of structured labels and temporal label dependencies. According to our experiments, the siLSTM provides a boost of 2.1% on THUMOS and 1.0% on MultiTHUMOS over the LSTM baseline as we can see from Table 6.1. In Figure 6.4, we present class-specific APs comparison between the siLSTM and LSTM. In Figure 6.5, we present the detailed overall improvement achieved, using siLSTM over the single-frame CNN baseline.

The significant difference obtained between our LSTM baseline and the LSTM results from [57] impede a direct and fair comparison with MultiLSTM. However, analyzing the relative gain between siLSTM and MultiLSTM over the corresponding baselines indicate that siLSTM is competitive with MultiLSTM. As presented before the relative gain provided by siLSTM is 2.1% and 1.0% against 2.0% and 1.6% for MultiLSTM on THUMOS and MultiTHUMOS, respectively. It is worth to mention that siLSTM and MultiLSTM explore different ideas for improving over LSTM, which are not mutually exclusive thus, they can possibly be combined for obtaining a potentially stronger model.

Figure 6.6 shows a timeline comparison for four different videos in their entirety, comparing a set of key models covered in this work (i.e. single-frame CNN and SINN, LSTM, siLSTM). The series points for prediction in this plot were obtained by thresholding the prediction at 0.5, similar procedure was done in [57].

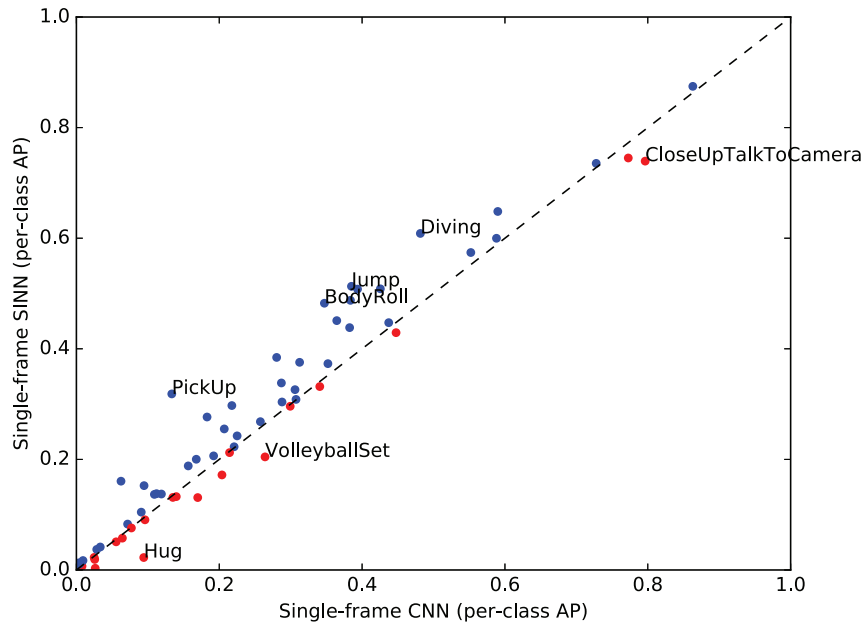


Figure 6.2: Per-class APs for Single-frame SINN versus Single-frame CNN (VGG).




Figure 6.3: Per-class APs for LSTM versus Single-frame CNN (VGG).

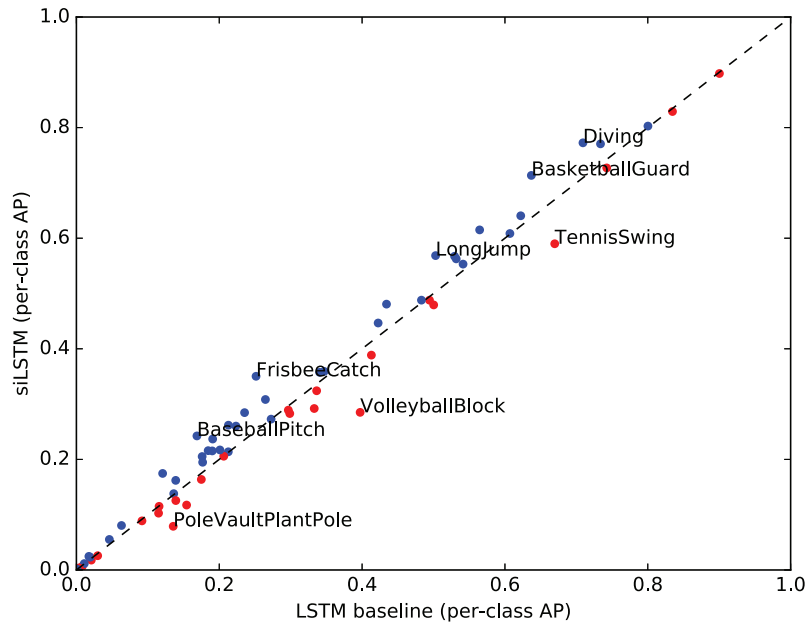


Figure 6.4: Per-class APs for siLSTM versus LSTM.




Figure 6.5: Per-class APs for siLSTM versus Single-frame CNN (VGG).

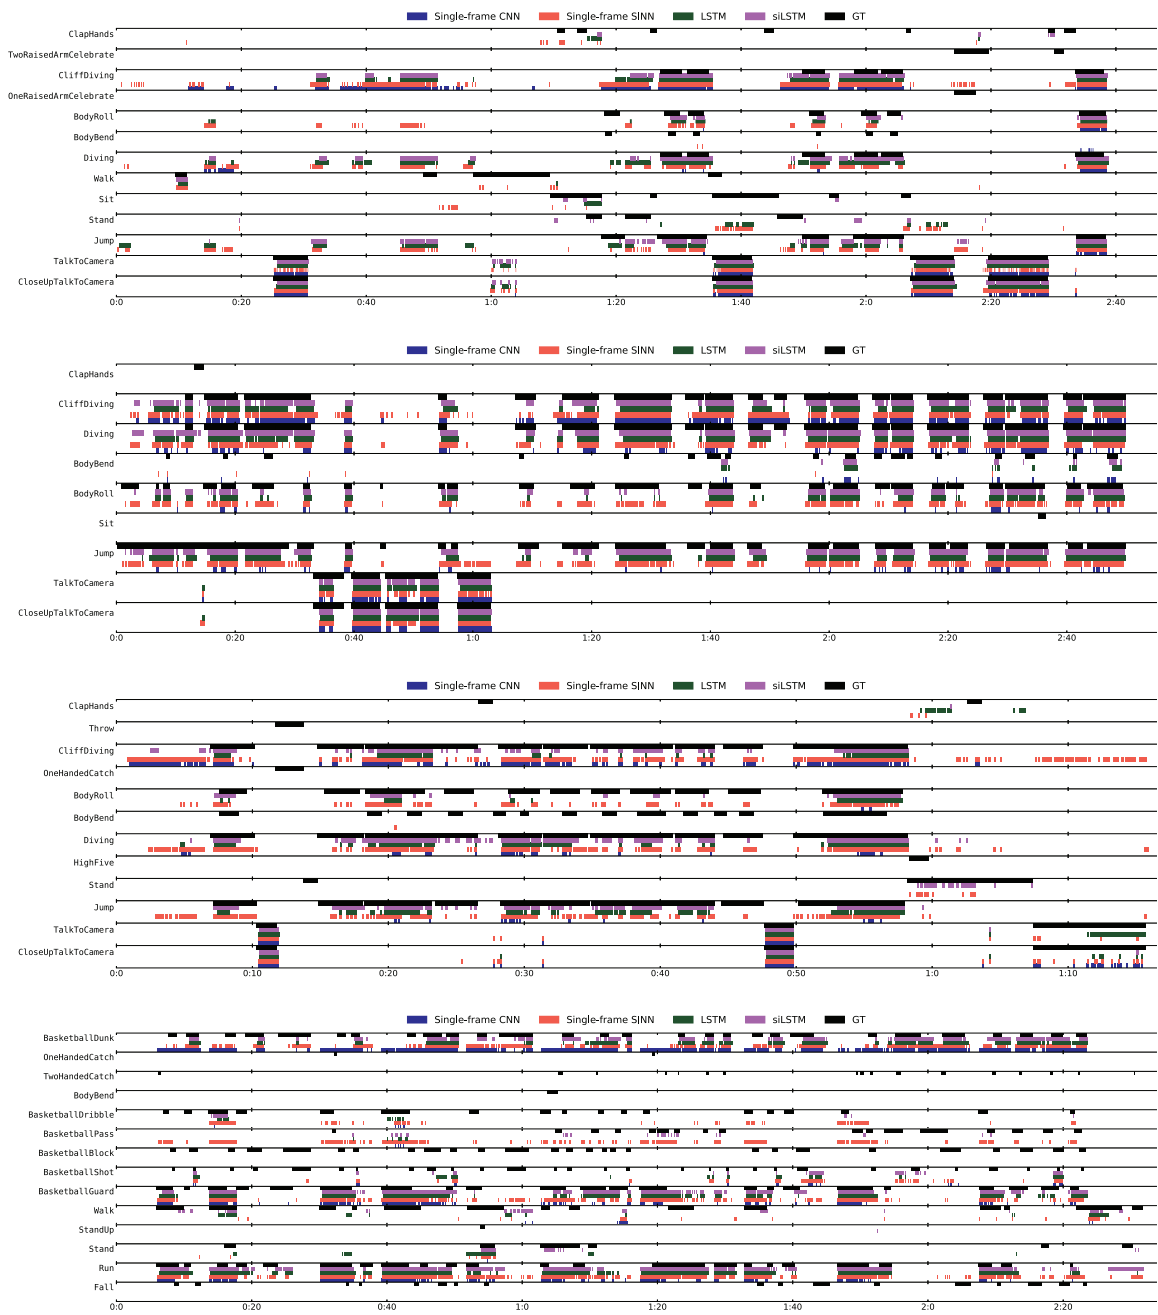


Figure 6.6: Timeline comparison for single-frame models (CNN and SINN) and sequential models (LSTM and siLSTM) in multi-label action detection for four different videos. The vertical axis represents the ground-truth labels appearing at least once in the target video and the horizontal axis corresponds to the duration of the target video in minutes. (Best view in color)

Chapter 7

Conclusion

In this thesis we bind together previous successes in performing inference in structured graphs (i.e. BINN and SINN), present new results for partially observed prediction on both releases of YouTube-8M and propose an extension for propagating concepts through time, based on the LSTM formulation. We consistently show significant improvements of performance in multi-label image classification, multi-label video classification and action detection tasks across a number of public datasets, achieving considerable gains against baselines and existing approaches at the cost of an additional computational overhead by performing structured inference.

From our experiments, the structured inference performed by BINN and SINN can lead to impressive boosts in accuracy. The results obtained for multi-label image and video classification are strongly positive and even without explicitly defining the graph structure as prior knowledge in the SINN model we were able to obtain substantial improvements on the YouTube-8M benchmarks. Additionally, the results presented using partially observed labels indicate that our method processes partial observations more effectively than the baselines and is able to predict missing labels with higher confidence.

The proposed action detection method also provides some improvements however, we believe that the lack of a more atomic structure of actions, person-centric spatial annotations and higher confidence structured annotations pose challenges for the learning of less noisy relations in label space.

Bibliography

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *CoRR*, 2016.
- [2] Moez Baccouche, Franck Mamalet, Christian Wolf, Christophe Garcia, and Atilla Baskurt. Sequential deep learning for human action recognition. In Albert Ali Salah and Bruno Lepri, editors, *HBU*, 2011.
- [3] Tat-Seng Chua, Jinhui Tang, Richang Hong, Haojie Li, Zhiping Luo, and Yantao Zheng. Nus-wide: A real-world web image database from national university of singapore. In *CIVR*, 2009.
- [4] Jia Deng, Nan Ding, Yangqing Jia, Andrea Frome, Kevin Murphy, Samy Bengio, Yuan Li, Hartmut Neven, and Hartwig Adam. Large-scale object classification using label relation graphs. In *ECCV*, 2014.
- [5] Jia Deng, Sanjeev Satheesh, Alexander C. Berg, and Fei Li. Fast and balanced: Efficient label tree learning for large scale object recognition. In *NIPS*. 2011.
- [6] N. Ding, J. Deng, K. P. Murphy, and H. Neven. Probabilistic label relation graphs with ising models. In *ICCV*, 2015.
- [7] J. Donahue, L. A. Hendricks, M. Rohrbach, S. Venugopalan, S. Guadarrama, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(4):677–691, 2017.
- [8] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014.
- [9] Yunchao Gong, Yangqing Jia, Thomas Leung, Alexander Toshev, and Sergey Ioffe. Deep convolutional ranking for multilabel image annotation. *CoRR*, 2013.
- [10] Kristen Grauman, Fei Sha, and Sung Ju Hwang. Learning a tree of metrics with disjoint visual features. In *NIPS*. 2011.
- [11] Alex Graves and Jürgen Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. In *NIPS*. 2009.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

- [13] Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *ICASSP*. 2017.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation (NC)*, 9(8):1735–1780, November 1997.
- [15] Hexiang Hu, Guang-Tong Zhou, Zhiwei Deng, Zicheng Liao, and Greg Mori. Learning structured inference neural networks with label relations. In *CVPR*, 2016.
- [16] Sung J Hwang, Kristen Grauman, and Fei Sha. Semantic kernel forests from multiple taxonomies. In *NIPS*, 2012.
- [17] Sung Ju Hwang and Leonid Sigal. A unified semantic embedding: Relating taxonomies and attributes. In *NIPS*. 2014.
- [18] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 35(1):221–231, 2013.
- [19] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM MM*, 2014.
- [20] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes, 2014.
- [21] Justin Johnson, Lamberto Ballan, and Li Fei-Fei. Love thy neighbors: Image annotation by exploiting image metadata. In *ICCV*, 2015.
- [22] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(4):664–676, 2017.
- [23] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [24] Y. Ke, R. Sukthankar, and M. Hebert. Event detection in crowded videos. In *ICCV*, 2007.
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, 2014.
- [26] Pushmeet Kohli, Jonathan Rihan, Matthieu Bray, and Philip H. S. Torr. Simultaneous segmentation and pose estimation of humans using dynamic graph cuts. *International Journal of Computer Vision (IJCV)*, 79(3):285–298, 2008.
- [27] Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. What are you talking about? text-to-image coreference. In *CVPR*, 2014.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

- [29] M. P. Kumar, P. H. S. Ton, and A. Zisserman. Obj cut. In *CVPR*, 2005.
- [30] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 36(3):453–465, 2014.
- [31] Bastian Leibe, Ales Leonardis, and Bernt Schiele. Combined object categorization and segmentation with an implicit shape model. In *ECCV Workshop*, 2004.
- [32] Haizhou Li, Helen M. Meng, Bin Ma, Engsiong Chng, and Lei Xie, editors. *Interspeech*, 2014.
- [33] F. Lv and R. Nevatia. Single view human action recognition using key pose matching and viterbi path searching. In *CVPR*, 2007.
- [34] S. Ma, L. Sigal, and S. Sclaroff. Learning activity progression in lstms for activity detection and early detection. In *CVPR*, 2016.
- [35] Julian J. McAuley and Jure Leskovec. Image labeling on a network: Using social-network metadata for image classification. In *ECCV*, 2012.
- [36] George A. Miller. Wordnet: A lexical database for english. *Communications of the ACM (CACM)*, 38(11):39–41, 1995.
- [37] Alberto Montes, Amaia Salvador, Santiago Pascual, and Xavier Giro-i Nieto. Temporal activity detection in untrimmed videos with recurrent neural networks. In *NIPS Workshop*, 2016.
- [38] Nelson Nauata, Jonathan Smith, and Greg Mori. Hierarchical label inference for video classification. *CoRR*, 2017.
- [39] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.
- [40] Siddharth V Patwardhan. *Incorporating Dictionary and Corpus Information into a Context Vector Measure of Semantic Relatedness*. Master’s thesis, University of Minnesota, 2003.
- [41] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [42] Pierre Sermanet, David Eigen, Xiang Zhang, Michael Mathieu, Rob Fergus, and Yann Lecun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *ICLR*, 2014.
- [43] Qinfeng Shi, Li Cheng, Li Wang, and Alex Smola. Human action segmentation and recognition using discriminative semi-markov models. *International Journal of Computer Vision (IJCV)*, 93(1):22–32, 2011.

- [44] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, 2014.
- [45] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, 2016.
- [46] K. Soomro, A. Roshan Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. In *CRCV-TR-12-01*, 2012.
- [47] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [48] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.
- [49] Ben Taskar, Carlos Guestrin, and Daphne Koller. Max-margin markov networks. In *NIPS*, 2003.
- [50] Ioannis Tsochantaridis, Thorsten Joachims, Thomas Hofmann, and Yasemin Altun. Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research (JMLR)*, 6:1453–1484, 2005.
- [51] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [52] Heng Wang, A. Klaser, C. Schmid, and Cheng-Lin Liu. Action recognition by dense trajectories. In *CVPR*, 2011.
- [53] B. Wu and R. Nevatia. Simultaneous object detection and segmentation by boosting local shape feature based classifier. In *CVPR*, 2007.
- [54] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010.
- [55] Jianxiong Xiao, Krista A. Ehinger, James Hays, Antonio Torralba, and Aude Oliva. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision (IJCV)*, 119(1):3–22, 2016.
- [56] J. Yamato, J. Ohya, and K. Ishii. Recognizing human action in time-sequential images using hidden markov model. In *CVPR*, 1992.
- [57] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, 2017.
- [58] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. In *NIPS*. 2014.