

3D Human Activity Recognition with Reconfigurable Convolutional Neural Networks

Keze Wang¹, Xiaolong Wang¹, Liang Lin^{1*}, Meng Wang², Wangmeng Zuo³

¹Sun Yat-Sen University, Guangzhou 510006, China

²School of Computer Science and Information Engineering, Hefei University of Technology

³School of Computer Science and Technology, Harbin Institute of Technology

linliang@ieee.org; {kezewang,dragonwxl123,eric.mengwang,cswmzuo}@gmail.com

ABSTRACT

Human activity understanding with 3D/depth sensors has received increasing attention in multimedia processing and interactions. This work targets on developing a novel deep model for automatic activity recognition from RGB-D videos. We represent each human activity as an ensemble of cubic-like video segments, and learn to discover the temporal structures for a category of activities, i.e. how the activities to be decomposed in terms of classification. Our model can be regarded as a structured deep architecture, as it extends the convolutional neural networks (CNNs) by incorporating structure alternatives. Specifically, we build the network consisting of 3D convolutions and max-pooling operators over the video segments, and introduce the latent variables in each convolutional layer manipulating the activation of neurons. Our model thus advances existing approaches in two aspects: (i) it acts directly on the raw inputs (grayscale-depth data) to conduct recognition instead of relying on hand-crafted features, and (ii) the model structure can be dynamically adjusted accounting for the temporal variations of human activities, i.e. the network configuration is allowed to be partially activated during inference. For model training, we propose an EM-type optimization method that iteratively (i) discovers the latent structure by determining the decomposed actions for each training example, and (ii) learns the network parameters by using the back-propagation algorithm. Our approach is validated in challenging scenarios, and outperforms state-of-the-art methods. A large human activity database of RGB-D videos is presented in addition.

*Corresponding author is Liang Lin. This work was supported by the Hi-Tech Research and Development Program of China (no.2013AA013801), Guangdong Natural Science Foundation (no.S2013050014548), Program of Guangzhou Zhujiang Star of Science and Technology (no.2013J2200067), Special Project on Integration of Industry, Education and Research of Guangdong (no.2012B091100148), and Fundamental Research Funds for the Central Universities (no.13lgjc26).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM'14, November 3–7 2014, Orlando, FL, USA.

Copyright 2014 ACM 978-1-4503-3063-3/14/11 ...\$15.00.

<http://dx.doi.org/10.1145/2647868.2654912>.

Categories and Subject Descriptors

I.4 [Computing Methodologies]: Image Processing and Computer Vision

General Terms

Algorithms; Experimentation; Performance

Keywords

Video Parsing; 3D Activity; Deep Learning; Structured Model

1. INTRODUCTION

In the research of multimedia, there is a particular interest in the last decades on developing intelligent systems of human activity understanding with different application backgrounds, e.g. intelligent surveillance, robotics and video content search. Recently developed 3D/depth sensors have opened up new opportunities with enormous commercial values, which provide more rich information compared with the traditional cameras (e.g., three-dimensional structure information of scenes and subjects/objects). Building upon these technologies, human poses can be easily access to, and modeling complicated human activities becomes facilitated.

This paper focuses on recognizing complex human activities from RGB-D videos that are captured by a Microsoft Kinect camera. There exist two main difficulties despite the additionally provided depth information:

- Representing complex human appearance and motion information. Due to diverse poses/views of individuals, it is usually hard to retrieve accurate body information with motions. The depth maps are often unavoidably contaminated [20], and may become unstable due to sensor noises or self-occlusions of bodies.
- Capturing large temporal variations of human activities. An activity can be considered as a sequence of actions occurred over time [3]. For instance, the activity of “microwaving food” can be decomposed into several actions such as picking up food, walking and operating microwave. The variance of a category of human activities can be hence very large, as it is uncertain how the activities to be decomposed in the temporal domain. Figure 1 shows two activities belonging to the same category, where the temporal lengths of decomposed actions are variant by different subjects.

Most of previous methods recognize 3D human activities by training discriminative or generative classifiers based on carefully designed features [36, 20, 35, 30]. These approaches

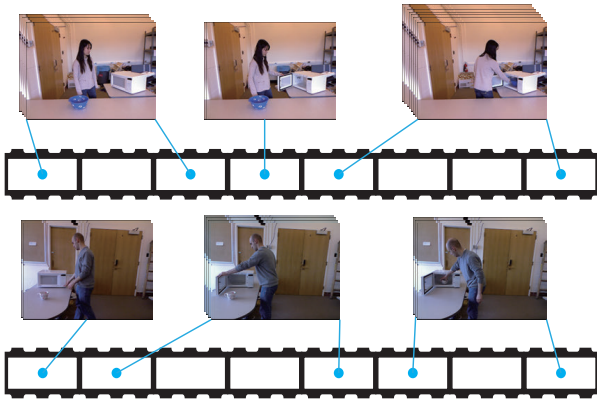


Figure 1: Two activities of the same category. We consider one activity as a sequence of actions occurred over time, and temporal compositions of actions are diverse by different subjects.

often require sufficient domain knowledge and heavy burden of feature engineering, which could limit their applications. Some compositional methods [33, 3] attempt to model activities by representing videos as sequences of fixed length temporal segments. However, they may have problems on handling complex activities composed by actions of diverse temporal durations, e.g., the example in Figure 1.

In this work, we develop an expressive configurable human activity model to address the above mentioned issues, absorbing the powers of two promising techniques of broad interests: deep learning [13, 7, 11, 8, 34, 18, 32] and reconfigurable part-based models [41, 1, 22, 16]. We represent one human activity as a sequence of separated actions, each of which is associated with a cubic-like video segment of unfixed length, as Figure 2 illustrates, and learn to discover the temporal structures for a category of human activities in terms of classification. In brief, our model is built upon the deep convolutional neural networks (CNNs) [13, 8], and we allow the network to be reconfigured to capture the varying temporal compositions of activities. We thus regard our model as a deep structured model, as it incorporates structure alternatives into a deep architecture. We consider following advantages of the deep structured model for 3D human activity recognition.

First, the deep architecture enables us to act directly on grayscale-depth data rather than relying on hand-crafted features. We build the layered network stacked up by convolutional layers, max-pooling operators and full connection layers, where the raw segmented videos are treated as inputs. We firstly apply the 3D convolutional kernel [8] over the bottom to extract features from both spatial and temporal domains, thereby encoding the motion information over adjacent frames. The 2D convolutions are then deployed upon to abstract the higher-level information. The convolutional layers for the segmented videos are computed independently to each other, in order to generate features for actions within the video segments. Afterwards, the convolution results coming from different segments are merged together into two full connection layers, giving rise to the activity classification.

Second, the structure of our model can be flexibly adjusted during inference, which is a key to improve the capability of

modeling complex patterns [41, 22, 14]. Specifically, in each convolutional layer, we impose the latent variables to manipulate the activation of neurons, so that the network can be partially enabled to explicitly handle large temporal variations of activities. For example, some of the neurons can be turned-off to adapt the different temporal durations of separated actions. During the inference for activity recognition, we aggregate the responses in each layer of network while searching for the optimal network configuration. It is worth mentioning that we can conduct the inference in a parallel manner using GPU (Graphic Processing Unit) programming, in order to counterbalance the extra computational demand.

Training the structured deep model is another innovation of this work, as it is required to simultaneously optimize parameters and latent structure in the deep architecture. Thus we propose an EM-type optimization method, namely Latent Structural Back Propagation (LSBP), which iterates with two steps: (i) Fixing the current model parameters, it performs activity classification while discovering the temporal composition (i.e. determining the separated actions) for each training example. (ii) Fixing the decompositions of input videos, it learns the parameters in each layer of the network using the back-propagation algorithm.

Moreover, collecting RGB-D data is relatively expensive in practice, while the amount of training data plays a critical role in deep feature learning [11]. Thus we propose to pre-train the network on the common RGB video data, taking advantage of existing databases of human activities. The trained parameters are then transferred into our model as the initializations.

The key contribution of this work is a novel deep structured model. To the best of our knowledge, it is original in literature to make the deep architecture reconfigurable to adaptively account for data variations. We demonstrate superior performances over state-of-the-art approaches in several challenging scenarios. In addition, we construct a new database of RGB-D data, which includes 1180 human activities of 20 categories.

This paper is organized as follows. Section 2 presents a review of related work. Then we present our deep structured model in Section 3, followed by a description of model learning algorithm in Section 4. Section 5 discusses the inference procedure. The experimental results and comparisons are exhibited in Section 6. Section 7 concludes this paper.

2. RELATED WORK

A batch of works on human action/activity understanding mainly focused on developing robust and descriptive features [35, 6, 20, 19, 40, 37, 24]. Xia and Aggarwal [35] extracted spatio-temporal interest points from depth videos (DSTIP) and developed a depth cuboid similarity feature (DCSF) to model human activities. Oreifej and Liu [20] proposed to capture spatio-temporal changes of activities by using a histogram of oriented 4D surface normals (HON4D). Most of these methods, however, overlooked detailed spatio-temporal structure information, and limited in periodic activities.

Several compositional approaches were studied for complex scenarios and achieved substantial progresses [33, 29, 39, 21, 23, 30, 4, 15], and they decomposed an activity into deformable parts and enriched the models with contextual information. For instance, Wang et al. [33] recognized

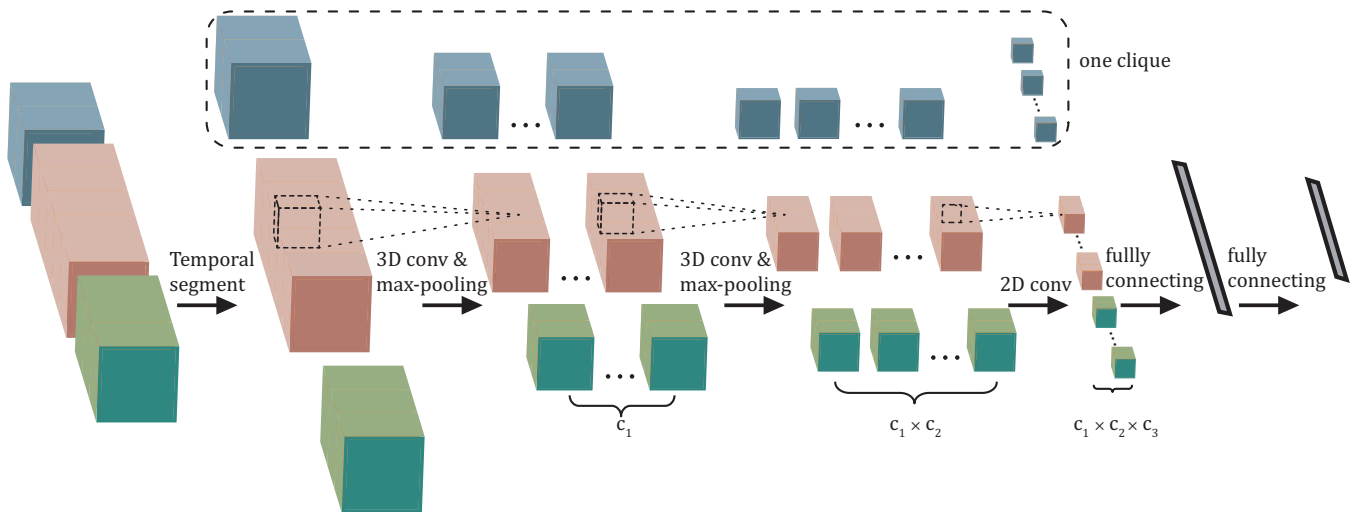


Figure 2: The architecture of our deep structured model. The network is stacked up by convolutional layers, max-pooling operators and full connection layers, where the raw segmented videos are treated as the input. A clique is defined as a subpart of the network stacked up for several layers, extracting features for one segmented video. Moreover, the architecture can be partially enabled to explicitly handle different temporal compositions of the activities.

human activities in common videos by training the hidden conditional random fields in a max-margin framework. For activity recognition in RGB-D data, Packer et al. [21] employed the latent structural SVM to train the model with part-based pose trajectories and object manipulations. An ensemble model of actionlets were studied in [30] to represent 3D human activities with a new feature called local occupancy pattern (LOP). To handle more complicated activities with large temporal variations, some powerful models [26, 31, 2] further discovered temporal structures of activities by localizing sequential actions. For example, Wang and Wu [31] proposed to solve the temporal alignment of actions by maximum margin temporal warping. Tang et al. [26] captured the latent temporal structures of 2D activities based on the variable-duration hidden Markov model. Koppala and Saxena [10] applied the Conditional Random Fields to model the sub-activities and affordances of the objects for 3D activity recognition.

Recently, the reconfigurable models were developed in the form of And-Or graphs [41, 22, 14, 28, 32], and yielded competitive performance in several challenging scenarios. The key idea of these approaches was to discover different ways of compositions by making the models reconfigured during learning and inference. Zhu and Mumford [41] first explored the And-Or graph models for image parsing. Pei et al. [22] then introduced the models for video event understanding, but their approach required elaborate annotations. Liang et al. [14] proposed to automatically train the reconfigurable action model by a non-convex formulation. However, the above mentioned models were built on hand-crafted features.

On the other hand, the past few years have seen a resurgence of research in the design of deep neural networks, and impressive progresses were made on learning image features from raw data [7, 27, 12, 11, 17, 5]. To address human action recognition from videos, Ji et al. [8] developed a novel deep architecture of convolutional networks, where they extracted features from both spatial and temporal dimensions. Amer

and Todorovic [1] applied Sum Product Networks (SPNs) to model human activities based on variable primitive actions. Our deep structured model can be viewed as an extension of these existing architectures, in which we make the network reconfigurable during learning and inference.

3. STRUCTURED DEEP MODEL

In this section, we firstly introduce the structure of our deep structured model, and then explain how it can handle large intra-class variance with the latent structure.

3.1 Spatio-temporal CNNs

Our deep model is presented as a spatio-temporal convolutional neural network, as shown in Figure 2. To model the complex human activities, it comprises of M network cliques, which jointly conduct the final output. We define a clique as a subpart of the network stacked up for several layers. In particular, each clique extracts features from one decomposed video segment associated to one separated sub-action from the complete activity, and an illustration is highlighted in Figure 2. Specifically, for each clique, two 3D convolutional layers are first built upon the raw input (i.e. grayscale and depth data), which consists with at most m video frames, and then followed by one 2D convolutional layer. Note that a max-pooling operator is applied on each 3D convolutional layer making our model robust to local body deformations and surrounding noises. Afterwards, the convolution results generated by different cliques are merged and concatenated into a long feature vector, upon which we build two full connection layers to associate with the activity labels. In the following, we introduce the detailed definitions for these components of our model.

3D Convolutional Layer. The 3D convolutional kernels in one clique are computed independently to those from different cliques. For example, the kernels belonging to the first clique will only be applied to perform convolutions on

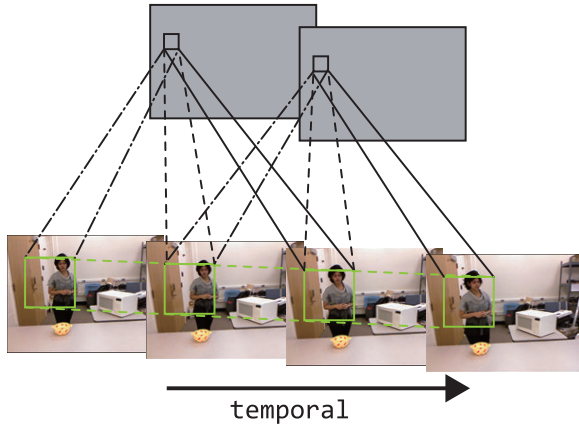


Figure 3: Illustration of the 3D convolution across both spatial and temporal domains. In this example, the temporal dimension of the 3D kernel is 3. That is to say, each feature map is obtained by performing 3D convolutions across 3 adjacent frames.

the first temporal segment of the activity. Motivated by [8], we perform the 3D convolutions spanning over both spatial and temporal dimensions of the input videos, and our model thereby captures both appearance and motion information for the observations. Suppose the width and the height of each frame are w and h , and the size of the 3D convolutional kernel is $w' \times h' \times m'$, where w' , h' , m' represents the width, height and temporal length, respectively. As Figure 3 illustrates, we can obtain a feature map via performing 3D convolutions across the s th to the $(s + m' - 1)$ th frames. The response for the position (x, y) in the feature map can be represented as,

$$v_{xys} = \tanh\left(b + \sum_{i=0}^{w'-1} \sum_{j=0}^{h'-1} \sum_{k=0}^{m'-1} \omega_{ijk} \cdot p_{(x+i)(y+j)(s+k)}\right), \quad (1)$$

where $p_{(x+i)(y+j)(s+k)}$ is the input pixel value at position $(x + i, y + j)$ in the $(s + k)$ th frame, ω_{ijk} is the parameter for the convolutional kernel, and b is the bias for the feature map. Thus we can obtain $m - m' + 1$ feature maps, each with size of $(w - w' + 1, h - h' + 1)$. Note that one convolutional kernel only extracts one kind of feature. Thus we employ several kernels to generate different kinds of feature in each convolutional layer. For each model clique, we define that the number of 3D convolutional kernels in the first and second layers as c_1 and c_2 .

After the first 3D convolutions, we obtain c_1 sets of $m - m' + 1$ feature maps. For each set of feature maps, we further perform 3D convolutions on it, and generate another set of feature maps on a deeper layer. Note that as we employ c_2 kernels on the c_1 sets of feature maps, we can obtain $c_1 \times c_2$ sets of new feature maps in the next layer.

Max-pooling Operator. In our model, we apply a max-pooling operator after each 3D convolution result. This is a procedure widely applied to obtain deformation and shift invariance [11, 38]. Given a set of feature maps, the max-pooling operator performs subsampling on them, which leads to the same number of feature maps with lower spatial resolution. More specifically, if a 2×2 max-pooling operator is performed on a $a_1 \times a_2$ feature map, we collect the max

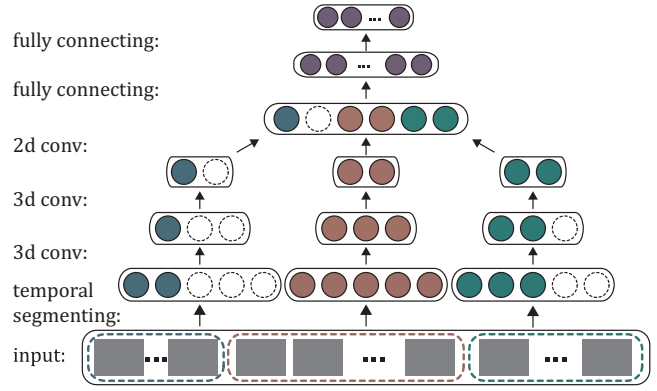


Figure 4: Illustration of our deep model incorporating the latent structure. Different model cliques are represented by different colors. As the temporal segmentation for an input video can be variant, different cliques might have different number of input frames. For the cliques whose input frame number less than m , part of the neurons in the network are inactivated, as represented by the dotted blank circles.

value in each 2×2 non-overlap regions to form a new feature map with size of $a_1/2 \times a_2/2$.

2D Convolutional Layer. After two layers of 3D convolution followed with max-pooling, each set of feature maps reduces to relatively a small temporal dimension. We then further apply 2D convolutional kernels to extract higher-level complex features from them. The 2D convolution can be viewed as a special case of 3D convolution by setting the temporal dimension of the 3D kernel to one, i.e. $m' = 1$. By performing 2D convolution on a set of feature maps, we can obtain the same number of feature maps in a new set. Suppose the number of 2D convolutional kernels is c_3 , by performing 2D convolutions on $c_1 \times c_2$ sets of feature maps, we can obtain $c_1 \times c_2 \times c_3$ sets of new feature maps.

Full Connection Layer. There are two full connection layers in our model. We first concatenate different sets of feature maps from the M model cliques into a long feature vector. Then each unit of this vector are further connected with all neurons in the first full connection layer, which are further fully connected with output neurons. Note that the number of the output neurons is K , which is the same as the number of categories of activities K , and each of the neurons represents the probability of an activity hypothesis. To normalize the probabilities of the output labels, we apply the softmax function on them,

$$\sigma(z_i) = \frac{\exp(z_i)}{\sum_{k=1}^K \exp(z_k)}, \quad (2)$$

where z_i is the i th value computed by multiplying the neuron values in the second last layer with the weights connected to the i th output neuron, and $\sigma(z_i)$ is the output probability. Note that $\sum_{i=1}^K \sigma(z_i) = 1$.

Input Data Details. In our experiment, we obtain the gray and depth image from the raw data, which are taken from the gray channel and the depth channel of each video frame. To perform convolutions, we duplicate the channels for the 3D convolutional kernels in the first layer. The convolution

results for these two channels are summed up together, thus the dimensions of the convolved feature maps remain the same. Note that our model can be generalized to apply on multi-channel video frames.

3.2 Reconfigurable Latent Structure

One key contribution of this paper is incorporating latent structure in deep model. Given different activity videos, the starting anchor frame and the number of input frames for each model clique can be variant. To illustrate it, we present a toy example in Figure 4, in which three model cliques represented by different colors are presented. Accordingly, the whole activity is decomposed into three action segments. The starting frame for each clique is flexible and the video is not evenly segmented. In this scheme, a clique is trained to handle the missing data by inactivating parts of the neurons. In other words, for the clique whose input anchor frame number is less than m , part of the neurons in the network are inactivated, as represented by the dotted blank circles in the first and third model cliques. Given input segments, classification on activities can be achieved with the deep model by performing forward propagation.

Formally, given a video sample, we define the index of starting anchor frames for M cliques as (s_1, \dots, s_M) and the corresponding number of input frames are (t_1, \dots, t_M) where $1 \leq t_i \leq m$. Thus, the latent variables for our model can be represented as $H = (s_1, \dots, s_M, t_1, \dots, t_M)$, which infers the video segments for each model clique. Given the input video X , latent variables H , and model parameters ω (including different layer parameters and biases), we represent the classification results obtained by our model as $F(X, \omega, H)$, which is a vector of the probabilities for each activity. For simplicity, we define the probability for the i th activity as $F_i(X, \omega, H)$.

Note that our model explicitly decomposes the input activity into sub-actions by incorporating the latent variables. But the sub-actions will be not directly co-related with semantic meaning. The different values of M (i.e. the number of sub-actions for one activity) will affect the final classification performance. When there is only one clique, the reconfigurability would be disregarded, and our model is simplified as the traditional CNN such as [8], which is compared in Sect. 6. We can empirically tune M , just like setting the number of parts for the deformable part-based model in object detection. The model with a small M could be less expressive to handle activity temporal variations, while a large M could lead to over-fitting for the high model complexity. In practice, we can roughly estimate M using traditional parameter tuning methods, e.g., the cross validation. We have shown the effectiveness of our method by setting $M = 4$ on the databases.

4. LEARNING

As our deep model incorporates latent structure, the standard back propagation algorithm [13] is not applicable to optimize the model parameters. Thus we propose an EM-type optimization algorithm, namely Latent Structural Back Propagation (LSBP), to learn our model. Due to the large number of our model parameters and the insufficiency of RGB-D data in human activities, we introduce a pre-training scheme to borrow the data strength from 2D data to optimize our model.

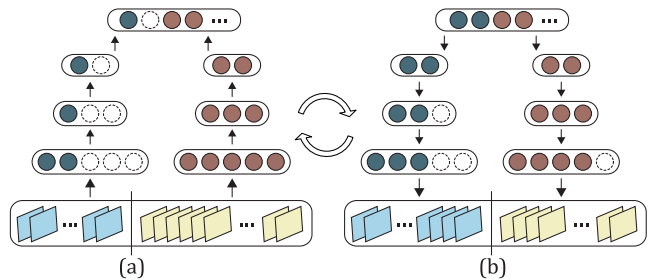


Figure 5: Illustration for the Latent Structural Back Propagation (LSBP). It is an iterative algorithm with two steps: (a) Given the current model parameters ω , estimate the latent variables H by adjusting the video segmentation; (b) Given the estimated latent variables H^* which infers a video segmentation proposal, perform back propagation to optimize model parameters ω . Note that the network neurons can be partially inactivated (as the dotted circles) according to the input segments.

4.1 Latent Structural Back Propagation

Given a video with different latent variables, the input segments for each model clique will be different. During learning, the model parameters ω and the latent variables H are required to be optimized at the same time. We propose an EM-type algorithm to iteratively optimize ω and H in two steps: (i) Given the model parameters ω we can compute the latent variables H (Figure 5.(a)); (ii) Given the input frames decided by H , we can perform back propagation to optimize the model parameters ω (Figure 5.(b)).

Suppose there are a set of N training samples $(X_1, y_1), \dots, (X_N, y_N)$, where X is the video, $y \in \{1, \dots, K\}$ represents the activity classes and K is the number of activity classes. For simplicity, we also define the set of latent variables for all samples as $\mathbf{H} = \{H_1, \dots, H_N\}$. We apply the logistic regression to define the cost function $J(\omega, \mathbf{H})$ in training, which is defined as,

$$J(\omega, \mathbf{H}) = -\frac{1}{N} \left(\sum_{i=1}^N \sum_{k=1}^K \mathbf{1}(y_i = k) \log F_k(X_i, \omega, H_i) + (1 - \mathbf{1}(y_i = k)) \log(1 - F_k(X_i, \omega, H_i)) \right) + \|\omega\|^2, \quad (3)$$

where $\mathbf{1}(\cdot) \in \{0, 1\}$ is the indicator function. The first two terms in Eq.(3) are the opposite of the log-likelihood and the last one is the regularization term. To minimize the cost $J(\omega, \mathbf{H})$, we optimize parameters ω and \mathbf{H} in a 2-steps iteration as below.

(i) Given the model parameters ω obtained from the last iteration, we can minimize Eq.(3) by maximizing the probability $F_{y_i}(X_i, \omega, H_i)$ for each sample (X_i, y_i) , which is achieved by finding the most appropriate latent variable H ,

$$H_i^* = \operatorname{argmax}_{H_i} F_{y_i}(X_i, \omega, H_i). \quad (4)$$

Recall that we apply softmax operator on the output results as Eq.(2), thus the maximization of $F_{y_i}(X_i, \omega, H_i)$ also depresses the probabilities of other labels $F_k(X_i, \omega, H_i), \forall k \neq y_i$. It leads to the increase of the log-likelihood and decrease the cost $J(\omega, \mathbf{H})$.

(ii) Given the latent variables $\mathbf{H} = \{H_1, \dots, H_N\}$ for each

sample, we can obtain the input sample segments for the deep structured model. After computing the cost $J(\omega, \mathbf{H})$ with the current inputs, we can obtain the gradients of $J(\omega, \mathbf{H})$ with respect to parameters ω . By performing the back propagation algorithm, we can further decrease the cost $J(\omega, \mathbf{H})$ and optimize the model parameters ω . Note that during back propagation, we apply stochastic gradient descent to update the parameters, and the update stops when it runs through all the training samples for one time.

The optimization algorithm iterates between these two steps until Eq.(3) converges.

4.2 Model Pre-training

To handle the insufficient RGB-D data for training and to boost the activity recognition results, we apply a pre-training scheme to initialize our model before optimizing the parameters with our LSBP algorithm on 3D data.

Given the large sum of 2D activity videos with labels, we train our deep model in a supervised manner. In this procedure, we first initialize the model parameters randomly, and each 2D video is evenly segmented to the number as the number of model cliques. To train on the 2D data, we directly apply back propagation instead of the proposed LSBP algorithm. It is mainly because of the following two reasons: (i) The initial model parameters are unstable, and it is not reliable to estimate latent variables H with them; (ii) The training efficiency can be improved without considering the latent variables.

After training on the 2D data, we apply the parameters of the convolutional layers to initialize our model. Note that we have gray and depth channels for each input frame in 3D data and only one gray channel for the 2D data. We thus duplicate the dimension of the 3D convolutional kernels in the first layer and initialize the parameters for the depth channel by the parameters for the gray channel. For the full connection layers, we set the parameters to random values. The reason is that we only need to borrow the feature learned in the 2D data, because the higher level information should be learned directly from the specific 3D activity dataset.

We summarize the whole learning procedure as Algorithm 1.

Algorithm 1 Learning Framework

Input:

The labelled 2D and 3D activity dataset.

Output:

Model parameters ω .

Initialization:

Pre-train the Spatial-Temporal CNN in the 2D dataset.

Learning on 3D dataset:

repeat

1. Estimate the latent variables \mathbf{H} by fixing model parameters ω .
2. Optimize ω given the input sample segments indicated by \mathbf{H} .

until $J(\omega, \mathbf{H})$ in Eq.(3) converges.

5. INFERENCE

The inference task is to recognize the category of the activity given a video X . Formally, we perform the standard procedure of brute search for the activity label y and the latent

variables H by maximizing the probability of $F_i(X, \omega, H)$,

$$(y^*, H^*) = \operatorname{argmax}_{(y, H)} F_y(X, \omega, H). \quad (5)$$

To do this, we search across all the labels $y(1 \leq y \leq K)$ and calculate the maximum probability $F_y(X, \omega, H)$ by optimizing H . Given the domain space of $H = (s_1, \dots, s_M, t_1, \dots, t_M)$, we constrain the input frame number for each model clique as $\tau \leq t_i \leq m$, and different video segments should not have overlaps (i.e., $s_i + t_i \leq s_{i+1}$). In all our experiments, we set the constant $\tau = 5$ during training and inference. We enumerate all the possibilities of H under these constraints, and calculate the corresponding probabilities $F_y(X, \omega, H)$ via forward propagations. By selecting the highest probability, we obtain the optimal $F_y(X, \omega, H^*)$. Though this optimization is a procedure of brute search, we can take advantage of parallel computation with GPU for that the forward propagations decided by different H are independent of each other. In our experiment, parallel computation via GPU highly accelerates the inference speed.

6. EXPERIMENT

We validate our approach on the CAD-120 activity dataset [9] and a larger dataset newly created by us, namely Office Activity (OA) dataset. Comparisons with state-of-arts and empirical analysis are presented in the experiments as well.

The CAD-120 dataset contains 120 RGB-D activity sequences of 10 categories, which is widely used in 3D human activity recognition. These activities were performed by four different subjects, and each activity was repeated three times by the same actor. The challenges on this dataset lie in the large variance in object appearance, human pose, and viewpoint. The proposed OA dataset is more comprehensive, covering the regular daily activities taken place in an office, and it is the largest activity dataset of RGB-D videos, which includes 1180 sequences. Three RGB-D sensors (i.e. Microsoft Kinect cameras) are utilized to capture data in three different viewpoints, and more than 10 actors are involved. The activities are captured in two different offices to increase the variability, where each actor performs the same activity twice for each viewpoint. It is worth mentioning that this dataset also includes the activities performed by two subjects with interactions. In particular, it is divided into two sub-sets, each of which contains 10 categories of activities: OA1 (activities by a single subject) and OA2 (activities by two subjects). The categories in OA1 are: *{answering-phones, arranging-files, eating, moving-objects, going-to-work, finding-objects, mopping, sleeping, taking-water, wandering}*, and in OA2, we have *{asking-and-away, called-away, carrying, chatting, delivering, eating-and-chatting, having-guest, showing, seeking-help, shaking-hands}*. The OA database is publicly accessible¹.

Several sampled frames and depth maps from the databases are exhibited in Figure 6.

Following [9], we apply four-fold cross validation procedure for testing in the CAD-120 dataset. That is, the model is trained on the activities of three subjects and tested on a new subject each time. The final outputs are averaged on the results of all four validations. For the OA dataset, we apply five-fold cross validation in the similar way.

¹<http://vision.sysu.edu.cn/projects/3d-activity/>



Figure 6: Activity examples from the testing databases. Several sampled frames and depth maps are presented. (a) CAD-120, (b) OA1, (c) OA2, respectively, show two activities of the same category selected from the three databases.

6.1 Implementation

Given the RGB-D videos, we first normalize them into the same temporal length as preprocessing. For each video, we extract 120 frames by removing the ones with similar appearances in gray-scale. Then we further obtain 30 anchor frames with a step size of 4 in the 120 frames. Suppose we index the 120 frames from 1 to 120, then the selected frames are indexed by 1, 5, 9, ..., 120. We apply these 30 anchor frames as the inputs for our model.

We scale the input frame to $w = 80$ and $h = 60$ in our experiments. The number of decomposed video segments (i.e. actions) is $M = 4$, and the length of the maximum number of input frames of each segment is $m = 9$. Recall that the actual input frames can be less than m , as we introduce the latent variables manipulating the temporal decomposition. The networks in each clique are of the same structures (e.g. kernels). For each clique, the number of 3D convolutional kernels in the first layer is $c_1 = 7$, and the size of the kernel is $9 \times 7 \times 3$, where each number represents the width, height and temporal length. In the second layer, the number of 3D kernels is $c_2 = 5$, and the size is $7 \times 7 \times 3$. We apply 3×3 max-pooling operator over the 3D convolutions. In the 2D convolutional layer, we have $c_3 = 4$ kernels with size of 6×4 . Hence we can obtain 700 feature maps with size 1×1 as the outputs for each network clique, and we merge the feature maps together into a vector of $700 \times 4 = 2800$ dimensions. Each unit in this vector is linked to 64 neurons in the next full connection layer. At last, the 64 neurons are fully connected to associate with the activity labels.

The experiments are executed on a desktop with an Intel

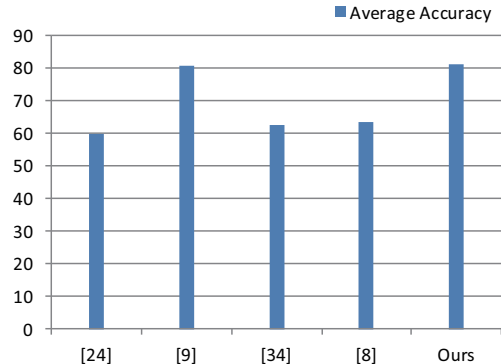


Figure 7: The average accuracy on the CAD-120 database.

i7 3.7GHz CPU, 8GB RAM and GTX TITAN GPU. For model learning, we set the learning rate as 0.002 for applying the stochastic gradient descent algorithm. The training times of one fold are 2 hours for CAD-120 (including 120 videos and occupying 0.2GB), 5 hours for OA1 (including 600 videos and occupying 1.0GB), and 5 hours for OA2 (including 580 videos and occupying 0.9GB) dataset, respectively. Each iteration of training costs similar time, and the convergence of our model over iterations is shown in Figure. 8 and 9. For inference, as we fully take advantage of parallel computation with GPU, it only takes around 0.4 seconds to complete recognition on a given video.

	Xia et al [35].	Ji et al [8].	Ours
arranging-objects	75.0%	68.3%	82.3%
cleaning-objects	68.3%	60.0%	79.7%
having-meal	41.7%	60.0%	71.0%
making-cereal	76.7%	77.6%	91.5%
microwaving-food	36.7%	71.7%	85.3%
picking-objects	75.0%	58.3%	97.2%
stacking-objects	75.0%	48.3%	61.0%
taking-food	83.3%	73.3%	93.5%
taking-medicine	58.3%	76.7%	96.8%
unstacking-objects	33.3%	36.7%	54.0%
Average accuracy	62.3%	63.1%	81.2%

Table 1: Accuracy of all categories on the CAD120 dataset.

	Xia et al [35]	Ji et al[8]	Ours
answering-phones	12.5%	40.0%	35.0%
arranging-files	59.7%	53.3%	84.4%
eating	40.3%	41.7%	65.5%
moving-objects	48.6%	51.7%	61.1%
going-to-work	34.7%	41.7%	92.2%
finding-objects	65.3%	36.7%	53.9%
mopping	63.9%	66.7%	72.2%
sleeping	25.0%	45%	43.9%
taking-water	58.3%	40.0%	51.7%
wandering	56.9%	50.0%	40.6%
Accuracy	46.5%	46.7%	60.1%

Table 2: Quantitative results on the OA1 dataset. Accuracy per activity category and average accuracy of all categories are reported.

	Xia et al [35]	Ji et al[8]	Ours
asking-and-away	12.5%	39.6%	25.3%
called-away	45.8%	44.8%	57.5%
carrying	66.7%	56.8%	53.5%
chatting	37.5%	17.2%	25.3%
delivering	20.1%	34.5%	32.8%
eating-and-chatting	50.0%	35.8%	69.5%
having-guest	37.5%	34.1%	43.7%
seeking-help	16.7%	44.8%	59.2%
shaking-hands	41.7%	32.8%	59.8%
showing	37.5%	29.3%	23.0%
Accuracy	36.6%	37.0%	45.0%

Table 3: Quantitative results on the OA2 dataset. Accuracy per activity category and average accuracy of all categories are reported.

	grayscale	depth	grayscale + depth
OA1	44.9%	57.2%	60.1%
OA2	41.6%	43.6%	45.0%

Table 4: Channel analysis on the three dataset. Average accuracy of all categories are reported.

6.2 Results and Comparisons

CAD-120 dataset. On this dataset, we adopt four state-of-the-art methods for comparison. As shown in Figure 7, our approach obtains the average accuracy of 81.2%, distinctly superior than results generated by other four competing methods, such as 59.7% [25], 80.6% [9], 62.3% [35] and 63.1% [8]. In Table 1, we report the detailed accuracies on all the categories, compared with the method based on hand-crafted feature engineering [35], and the deep architecture of convolutional neural networks [8]. Note that for different methods we train the models using the same data annotation, which only includes the activity labels on videos.

OA dataset. In this experiment, we apply our approach on the two sub-sets, respectively. Our deep structured model outperforms the state-of-the-art methods on average. On the OA1 set, our approach outperforms on 5 out of 10 categories and obtains the highest average accuracy of 60.1%, as Table 2 reports. On the OA2 set, the average accuracy of our method is 45.0% and we have 6 classes of activities achieve better results than the other two competing methods, as Table 3 reports. By reviewing the results, we find the failure cases probably caused by the lack of contextualized scene understanding. For example, understanding the activities of *taking-water* and *sleeping* actually requires to extra higher level information. We will consider it in future work. Moreover, the depth data is very useful. Since the testing is performed on the new subject, we observe that large appearance variances existed in grayscale data lead to worse performance. The depth data has much smaller variance, and does help to capture the motion information. Table 4 illustrates that depth data can boost the performance a lot, especially in OA1.

6.3 Empirical Analysis

For further evaluation, we conduct two following empirical analysis under different settings.

(I) To clarify significance of using the pre-training, we discard the parameters trained on 2D videos and learn the model directly on the RGB-D data. Then we compare the model with the original version by the test error rate, which is defined as one minus the classification accuracy. This testing is implemented on the OA1 dataset. In Figure 8, we visualize the test error rates with the increasing of iteration numbers during training. Each test error rate on a specific iteration number is calculated by applying the currently trained model. It is shown that the model using the pre-training converges after 25 iterations, while the other one without the pre-training requires 140 iterations. More importantly, the pre-training can effectively reduce the error rate in the testing, say 8% less than without the pre-training.

(II) We demonstrate the effectiveness of incorporating re-configurable structure in the deep architecture. That is, we can fix the latent variables in our model and train it using the standard back propagation algorithm. In this model without latent structure, each network clique receives m input frames that are evenly segmented from the video. In Figure 9, we visualize the test error rates with different iterations of the two models: structured v.s. non-structured, from the same initialization. We observe that the error rates of the structured model decrease faster and reach to the lower result,

compared with the non-structured model. This experiment is executed on the OA1 dataset.

7. CONCLUSIONS

This paper studies a novel deep structured model by incorporating model reconfigurability into layered convolutional neural networks. This model has been shown to handle well realistic challenges in 3D activity recognition, and it enables us to perform recognition from raw RGB-D data rather than relying on hand-crafted features. Moreover, we consider two aspects in future work. First, we can integrate high level semantic information into our model to deal with more complicated events with underlying intentions. Second, we plan to deploy our model into cloud computing platforms so that thin clients can expediently access the ability of 3D human activity understanding.

8. REFERENCES

- [1] M. R. Amer and S. Todorovic. Sum-product networks for modeling activities with stochastic structure. In *CVPR*, pages 1314–1321, 2012.
- [2] W. Brendel and S. Todorovic. Learning spatiotemporal graphs of human activities. In *ICCV*, pages 778–785, 2011.
- [3] J. M. Chaquet, E. J. Carmona, and A. Fernandez-Caballero. A survey of video datasets for human action and activity recognition. *Computer Vision and Image Understanding*, 117(6):633–659, 2013.
- [4] Z. Cheng, L. Qin, Q. Huang, S. Jiang, S. Yan, and Q. Tian. Human group activity analysis with fusion of motion and appearance information. In *ACM Multimedia*, pages 1401–1404, 2011.
- [5] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [6] R. Gupta, A. Y. Chia, D. Rajan, E. S. Ng, and E. H. Lung. Human activities recognition using depth images. In *ACM Multimedia*, pages 283–292, 2013.
- [7] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.
- [8] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(1):221–231, 2013.
- [9] H. S. Koppula, R. Gupta, and A. Saxena. Learning human activities and object affordances from rgb-d videos. *International Journal of Robotics Research (IJRR)*, 32(8):951–970, 2013.
- [10] H. S. Koppula and A. Saxena. Learning spatio-temporal structure from rgb-d videos for human activity detection and anticipation. In *ICML*, pages 792–800, 2013.
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [12] Q. V. Le, W. Y. Zou, S. Y. Yeung, and A. Y. Ng. Learning hierarchical invariant spatio-temporal

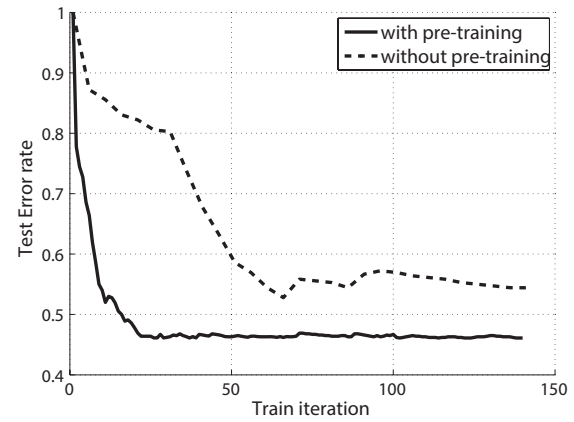


Figure 8: Test error rates with/without using the pre-training.

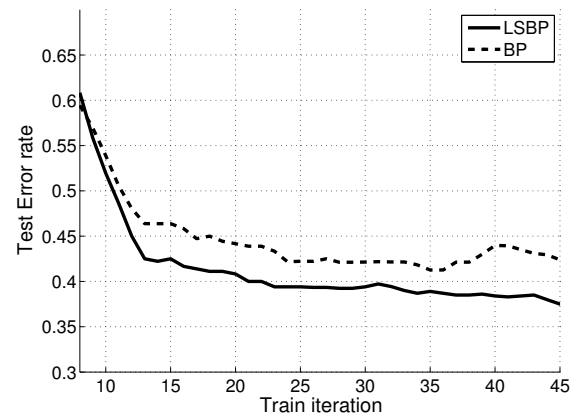


Figure 9: Test error rates with/without incorporating the latent structure in the deep model. The solid curve represents the structured deep model trained by the proposed LSBP optimization method, and the dashed curve represents the non-structured model using standard back propagation.

- features for action recognition with independent subspace analysis. In *CVPR*, pages 3361–3368, 2011.
- [13] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and e. a. L.D. Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*, 1990.
- [14] X. Liang, L. Lin, and L. Cao. Learning latent spatio-temporal compositional model for human action recognition. In *ACM Multimedia*, pages 263–272, 2013.
- [15] L. Lin, H. Gong, L. Li, and L. Wang. Semantic event representation and recognition using syntactic attribute graph grammar. *Pattern Recognition Letters*, 30(2):180–186, 2009.
- [16] L. Lin, T. Wu, J. Porway, and Z. Xu. A stochastic graph grammar for compositional object representation and recognition. *Pattern Recognition*, 42(7):1297–1307, 2009.
- [17] P. Luo, X. Wang, and X. Tang. A deep sum-product architecture for robust facial attributes analysis. In

- ICCV*, pages 2864–2871, 2013.
- [18] P. Luo, X. Wang, and X. Tang. Pedestrian parsing via deep decompositional neural network. In *ICCV*, pages 2648–2655, 2013.
- [19] B. Ni, Z. L. Y. Pei, L. Lin, and P. Moulin. Integrating multi-stage depth-induced contextual information for human action recognition and localization. In *International Conference and Workshops on Automatic Face and Gesture Recognition*, pages 1–8, 2013.
- [20] O. Oreifej and Z. Liu. Hon4d: Histogram of oriented 4d normals for activity recognition from depth sequences. In *CVPR*, pages 716–723, 2013.
- [21] B. Packer, K. Saenko, and D. Koller. A combined pose, object, and feature model for action understanding. In *CVPR*, pages 1378–1385, 2012.
- [22] M. Pei, Y. Jia, and S. Zhu. Parsing video events with goal inference and intent prediction. In *ICCV*, pages 487–494, 2011.
- [23] S. Sadanand and J. J. Corso. Action bank: A high-level representation of activity in video. In *CVPR*, pages 1234–1241, 2012.
- [24] P. Scovanner, S. Ali, and M. Shah. A 3-dimensional sift descriptor and its application to action recognition. In *ACM Multimedia*, pages 357–360, 2007.
- [25] J. Sung, C. Ponce, B. Selman, and A. Saxena. Unstructured human activity detection from rgb-d images. In *ICRA*, pages 842–849, 2012.
- [26] K. Tang, L. Fei-Fei, and D. Koller. Learning latent temporal structure for complex event detection. In *CVPR*, pages 1250–1257, 2012.
- [27] G. W. Taylor, R. Fergus, Y. LeCun, and C. Bregler. Convolutional learning of spatio-temporal features. In *ECCV*, pages 140–153, 2010.
- [28] K. Tu, M. Meng, M. W. Lee, T. Choi, and S. Zhu. Joint video and text parsing for understanding events and answering queries. *IEEE Transactions on Multimedia*, 21(2):42–70, 2014.
- [29] C. Wang, Y. Wang, and A. L. Yuille. An approach to pose-based action recognition. In *CVPR*, pages 915–922, 2013.
- [30] J. Wang, Z. Liu, Y. Wu, and J. Yuan. Mining actionlet ensemble for action recognition with depth cameras. In *CVPR*, pages 1290–1297, 2012.
- [31] J. Wang and Y. Wu. Learning maximum margin temporal warping for action recognition. In *ICCV*, pages 2688–2695, 2013.
- [32] X. Wang, L. Lin, L. Huang, and S. Yan. Incorporating structural alternatives and sharing into hierarchy for multiclass object recognition and detection. In *CVPR*, pages 3334–3341, 2013.
- [33] Y. Wang and G. Mori. Hidden part models for human action recognition: Probabilistic vs. max-margin. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(7):1310–1323, 2011.
- [34] P. Wu, S. Hoi, H. Xia, P. Zhao, D. Wang, and C. Miao. Online multimodal deep similarity learning with application to image retrieval. In *ACM Multimedia*, pages 153–162, 2013.
- [35] L. Xia and J. Aggarwal. Spatio-temporal depth cuboid similarity feature for activity recognition using depth camera. In *CVPR*, pages 2834–2841, 2013.
- [36] L. Xia, C. Chen, and J. K. Aggarwal. View invariant human action recognition using histograms of 3d joints. In *CVPRW*, pages 20–27, 2012.
- [37] X. Yang, C. Zhang, and Y. Tian. Recognizing actions using depth motion maps-based histograms of oriented gradients. In *ACM Multimedia*, pages 1057–1060, 2012.
- [38] K. Yu, Y. Lin, and J. Lafferty. Learning image representations from the pixel level via hierarchical sparse coding. In *CVPR*, pages 1713–1720, 2011.
- [39] X. Zhao, Y. Liu, and Y. Fu. Exploring discriminative pose sub-patterns for effective action classification. In *ACM Multimedia*, pages 273–282, 2013.
- [40] X. Zhou, X. Zhuang, S. Yan, S. F. Chang, M. H. Johnson, and T. S. Huang. Sift-bag kernel for video event analysis. In *ACM Multimedia*, pages 229–238, 2009.
- [41] S. Zhu and D. Mumford. A stochastic grammar of images. *Foundations and Trends in Computer Graphics and Vision*, 2(4):259–362, 2007.