Learning Latent Spatio-Temporal Compositional Model for Human Action Recognition

Xiaodan Liang Sun Yat-Sen University, China xdliang328@gmail.com Liang Lin* Sun Yat-Sen University, China linliang@ieee.org Liangliang Cao IBM Research, U.S. liangliang.cao@us.ibm.com

ABSTRACT

Action recognition is an important problem in multimedia understanding. This paper addresses this problem by building an expressive compositional action model. We model one action instance in the video with an ensemble of spatio-temporal compositions: a number of discrete temporal anchor frames, each of which is further decomposed to a layout of deformable parts. In this way, our model can identify a Spatio-Temporal And-Or Graph (STAOG) to represent the latent structure of actions e.g. triple jumping, swinging and high jumping. The STAOG model comprises four layers: (i) a batch of leaf-nodes in bottom for detecting various action parts within video patches; (ii) the or-nodes over bottom, *i.e.* switch variables to activate their children leaf-nodes for structural variability; (iii) the and-nodes within an anchor frame for verifying spatial composition; and (iv) the root-node at top for aggregating scores over temporal anchor frames. Moreover, the contextual interactions are defined between leaf-nodes in both spatial and temporal domains. For model training, we develop a novel weakly supervised learning algorithm which iteratively determines the structural configuration (e.g. the production of leaf-nodes associated with the or-nodes) along with the optimization of multi-layer parameters. By fully exploiting spatio-temporal compositions and interactions, our approach handles well large intra-class action variance (e.g. different views, individual appearances, spatio-temporal structures). The experimental results on the challenging databases demonstrate superior performance of our approach over other competing methods.

Categories and Subject Descriptors

I.5 [Computing Methodologies]: Pattern Recognition; I.4 [Computing Methodologies]: Image Processing and Computer Vision

General Terms

Algorithms, Experimentation, Performance

*Corresponding author is Liang Lin. This work was supported by the 863 Program of China (no.2012AA011504), the National Natural Science Foundation of China (no. 61173082), and the Special Project on the Integration of Industry, Education and Research of Guangdong Province (no. 2012B091100148), and the Guangdong Science and Technology Program (no. 2012B031500006).

Copyright 2013 ACM 0-89791-88-6/97/05 ...\$15.00.

Keywords

Video Understanding, Action Recognition, Structural Learning, And-Or Graph

1. INTRODUCTION

With the popularity of personal video cameras and multi-view video capturing devices, we are entering an era with rich amount of multimedia documents surrounding us. To interact with these videos, there have been increasing demands of understanding human activities in these videos. Although many research studies [45, 41, 31, 46, 22, 5, 29, 37, 32, 8] have been carried out to understand and retrieve large scale video contents, these works focus on high level semantics instead of describing human activities. There still exists a need to recognize fine-grained information for human activities, *e.g.* body poses and temporal motions.

This paper targets on the challenge of understanding spatial and temporal variances of video phenomenons. More specifically, we are considering the following difficulties:

- A human body is composed of multiple parts, and different parts are associated with different motions.
- Both the short term and long term motions may be fused by different background movement or camera motion, which brings the difficulties of accurately modeling the temporal characteristics for actions.
- In real world videos, human actions often happen with uncertainties: some happen along with occlusions, some are caused by view/pose variance, or due to diverse actor appearances and motions.

Due to these difficulties, it is crucial to reduce the spatio-temporal ambiguity when modeling the human actions. Most of the previous studies were built on simplified action models while overlooking the detailed spatio-temporal structure information. Of these works, a large amount of studies were based on spatio-temporal interest points [16, 42, 33]. Some researchers proposed to enrich the action model with the appearance information or context information [40, 34]. Some other researchers learned temporal structures for action recognition [28, 35]. However, none of these works provides an effective model which can unify spatial and temporal information to infer the structure of human motion.

We aims to develop an effective configurable model, namely the Spatio-Temporal And-Or Graph (STAOG) for action recognition, which addresses the problems mentioned above. Our idea is partially motivated by the image grammar model [47], which hierarchically decomposes an image pattern with mixed and-nodes and or-nodes, as well as modeling rich structural variations of parts.

The challenges of generalizing And-Or graphs for action recognition are two-folds. First, the traditional And-Or graphs are lim-

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.



Figure 1: An example of the Spatio-Temporal And-Or Graph model, where the and-nodes represent compositions in either time or space, the or-nodes indicate structural alternatives, and the leaf-nodes (at the bottom) correspond to local part detectors. The links between leaf-nodes represent spatio-temporal contextual interactions.

ited in modeling the hierarchical configuration of spatio-temporal information. Because actions in videos are often more complicated than images, we need more powerful models for video problems. Second, videos require more efficient models that can be effectively learned from large amount of video information without elaborate supervision and initialization.

To handle the first challenge, our STAOG model extends the traditional deformable graphical models by introducing switch variables in hierarchy, *i.e.* or-nodes that explicitly specify structural reconfiguration. Both spatial and temporal interactions between action parts can be simultaneously learned. One action in the video can be treated as an ensemble of spatio-temporal compositions: a number of discrete temporal anchor frames, each of which is further decomposed into a layout of deformable parts. An example of the proposed STAOG model is illustrated in Fig. 1. There are four layers in our model from bottom to top:

(1) The *leaf-nodes* in the bottom layer represent a batch of local classifiers for detecting various action parts in every anchor frame, denoted by the solid circles in Fig. 1. During detection, location displacement is allowed for each leaf-node to tackle the part deformation.

(2) The *or-nodes* over the bottom are "switch" variables specifying the activation of their children leaf-nodes, denoted by the dashed circles. Each or-node is used to specify an appropriate selection from candidate action parts detected by the associating children leaf-nodes. In this way, by explicitly switching selections over leaf-nodes, the or-nodes make our model reconfigurable during the inference of detection, which is the key to handle large action variabilities.

(3) The *and-nodes* verify the holistic appearance of action within the anchor frame, (the rectangles in layer 2), and we thus consider it as the spatial and-node. It includes two aspects: (i) a global classifier with bag-of-features, and (ii) aggregated scores from its children or-nodes.

(4) The root-node in the top can be viewed as an and-node in

time, (the rectangle in top). Its definition is similar to the spatial and-node: (i) a classifier with global features in observed frames, (ii) aggregated scores over candidate temporal anchor frames, plus penalty for anchor frame displacements.

(5) The *spatio-temporal contextual interactions*, *e.g.* the curves (graph edges) among leaf-nodes in Fig. 1, are defined based on informative contextual pairwise relations in either spatial or temporal domain. Note that the collaborative edges are imposed between leaf-nodes that are associated with different action parts. Their effectiveness will be particularly demonstrated in the experiments.

To overcome the second challenge, we present a novel weakly supervised learning algorithm for model training, inspired by the non-convex optimization techniques [43, 38]. This algorithm trains the model in a dynamic manner: the model structure (*e.g.* the configuration of leaf-nodes and or-nodes) is iteratively generated and reconfigured on the training data, with optimizing the multi-layer parameters. The other structure attributes (*e.g.* the activation of leaf-nodes, and temporal deformation of anchor frames) are modeled with the latent variables and optimized simultaneously.

In the testing stage, we present an algorithm of cascaded search and verification for recognizing actions with the trained STAOG model. We first generate a set of hypotheses in both spatial and temporal compositions. (i) Spatial testing via the and-nodes. Within the input frame, all candidate action parts are found by leafnodes and several possible configurations (*i.e.* spatial compositions of action parts) are produced with different specifications via the or-nodes. These configurations are also weighted via the and-node. (ii) Temporal testing via the root-node. The scores proposed by the and-nodes are aggregated via the root-node for a possible temporal composition. Several possible configurations are then produced as hypotheses represented by different latent variables. Finally, each hypothesis is globally verified with the spatial and temporal edges in the model.

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

2. RELATED WORK

Traditional works for action recognition focused on developing informative features, such as spatio-temporal descriptors [16, 7, 46],3D Gradient [15], 3D SIFT descriptors [29] and motion features [14, 18], and the action classifier can be trained with labeled data. Most of these methods, however, are limited to periodic actions with clean background, such as running and jogging.

To address complex actions with cluttered background, several compositional or expressive models were proposed and achieve very impressive results [23, 5, 27, 39, 13, 6]. For example, Wang et al. [39] modeled the human action by a flexible constellation of parts conditioned on image observations and learned the parameters of an HCRF model in a max-margin framework, motivated by the recent progresses in object recognition and detection, e.g. the deformable part model by Felzenszwalb et al. [10]. Yao et al. [40] proposed to generate spatio-template action templates with the information projection method. Sadanand et al [27] adopted highlevel representations with a bank of individual action detectors. However, actions in video often involve much more information in both spatial and temporal domain, compared with image-based object recognition, and most of these studies do not explicitly localize parts of actions (actors) due to the computational burden. Moreover, structural configurations of these models are usually fixed,

including a fixed number of part detectors as well as the predefined composition.

One unique characteristic of human action recognition problem lies in the temporal structure. A lot of works were proposed to build temporal structure models [9, 24, 35, 4] based on discriminative and interesting motion segments of the video. Raptis et al. [25] extracted clusters of trajectories and proposed a graphical model to incorporate constraints for individual and group events. Albanese et al. [1] represented temporal relations of activities using the probabilistic Petri Nets and integrated high-level reasoning approaches. Different from these approaches, we do not treat the whole temporal frames as units.Instead, we model temporal structure based on action parts with explicit relations and presents a solution to find both spatial and temporal configurations for dynamic activities.

Recently, the And-Or graph models [47] have been discussed for several vision tasks such as object recognition [20] and shape modeling [19]. These works mainly focused on images instead of videos and do not take the temporal dynamic structure into account. The very recent work by Amer et al. [3] proposed to recognize activities with the spatio-temporal And-Or graph model, but they over-simplified the model training by manually fixing the model structure (*i.e.* the layout of graph nodes).

It is worth mentioning that this paper learn the spatio-temporal graph without using any extra annotations or scripts. Research works which utilize rich annotation for event parsing and interpretation are beyond the scope of this work. In contrast, Marszalek et al. [23] explored the action contexts of natural dynamic scenes with movie scripts. Gupta et al. [11] proposed to learn a visually grounded storyline model from annotated videos, and Pei et al. [30] studied the event grammar model for daily activities based on a predefined set of unary and binary relations. Extra annotations are required for these studies.

3. SPATIO-TEMPORAL AND-OR GRAPH

The STAOG model is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the four types of nodes and \mathcal{E} the graph edges as Fig. 1. The root node in top verifies the temporal composition, which aggregates scores over anchor frames. Each and-node represents a temporal anchor frame for verifying spatial composition. The or-nodes are derived from each and-node, which are "switch" variables for specifying the activation of their children leaf-nodes. The number of leaf-nodes for each or-node is dynamically learned with an upper limit number m. For simplicity, we use t = 1, ..., T to index all and-nodes in the whole STAOG model, i = 1, ..., Z for ornodes and j = 1, ..., n for leaf-nodes. We also index the child or-node of and-node A_t as $i \in ch(t)$, and index the child leaf-node of or-node U_i as $j \in ch(i)$. The spatio-temporal graph edges(*i.e.* interactions) are defined between the leaf-nodes associated with different or-nodes. In this section, we describe two factors in detail: the spatio-temporal compositions, and the contextual interactions in both spatial and temporal domains.

3.1 Spatio-Temporal Compositions

We employ Laptev's 3-D corner detector [16] to detect interest points in video sequences, and each interest point is described by HoG (histogram of gradient) and HoF (histogram of optical flow) [17]. Furthermore, we generate a dictionary of spatio-temporal interest points' descriptors, clustered by the k-means method in training stage. Given a video sequence X, we first equally divide it into T temporal segments. The center frame in each video segment is chosen as an initial anchor frame. Each anchor frame is further decomposed into a number of action parts. In our method, we define the action parts based on the video patch representation,



Figure 2: Illustration of spatial compositions. (a) The black boxes denote the initial positions of action parts. (b) The parts are exhibited which are associated with a set of leaf-nodes. Each (d_x, d_y) indicates the location displacement determined by the model. (c) The activated leaf-nodes are highlighted by red and spatial contextual interactions are defined between the pairwise spatial adjacent leaf nodes.

i.e. 3-D volumes spanning ρ consecutive frames. Thus, for each anchor frame I_t , we observe a sequence of frames centered at I_t , and the sequence denoted as Λ_t is treated as the input for anchor frame processing.

Leaf-node: Each leaf-node L_j represents a local classifier for detecting action parts within video patches and includes two terms: an appearance feature ϕ^l and the spatial displacement feature ϕ^s . Within an anchor frame I_t , the features of action parts are described by the BoW histogram based on the generated dictionary. Assume the action part detected by L_j is localized at position $p_j = (p_j^x, p_j^y)$, then $\phi^l(\Lambda_t, p_j)$ is denoted as the appearance feature. During detection, the locations of action parts are allowed to be perturbed to tackle the spatial deformation. We incorporate the spatial displacement $\phi^s(q_j, p_j) = (d_x, d_y)$ for each action part, which can be computed by maximizing the response of L_j during inference; q_j , representing the initialized position of L_j , is set according to the center-point of the frame. Thus the spatial displacement is defined as $\phi^s(q_j, p_j) = (d_x, d_y)$, where (d_x, d_y) is the displacement. The response of L_j is defined as,

$$R_j^l(\Lambda_t, p_j) = \omega_j^l \cdot \phi^l(\Lambda_t, p_j) - \omega_j^s \cdot \phi^s(q_j, p_j), \qquad (1)$$

where ω_j^l is the parameter for the appearance feature and ω_j^s corresponds to the spatial deformation parameter.

Or-node: Each or-node U_i is proposed to specify an appropriate candidate from its children leaf-nodes. For each leaf-node L_j of U_i 's children, the indicator variable $v_j \in \{0, 1\}$ represents whether it is activated or not and each or-node only selects one leaf-node. Briefly, we utilize the indicator vector \mathbf{v}_i for the or-node U_i and each element of \mathbf{v}_i is an indicator variable v_j of the leaf-node L_j . Intuitively, the significant intra-class variance caused by views, background clutters or actors can be captured by different spatial configurations that are determined with the or-nodes. The response of the or-node U_i is defined as,

$$R_i^u(\Lambda_t, \mathbf{v}_i) = \sum_{j \in ch(i)} R_j^l(\Lambda_t, p_j) \cdot v_j,$$
(2)



Figure 3: Illustration of the contextual relations for defining spatial edges in the STAOG model. We define the edges between spatial adjacent leaf-nodes with 8 relations according to their spatial layout: above, below, left, right, near, far, clockwise and anti-clockwise.

And-node: Each and-node A_t verifies the holistic appearance of action for the anchor frame I_t , and spatial composition of the or-nodes in its children. We define the configuration vector V_t for all leaf-nodes within the anchor frame, which includes all indicator vectors \mathbf{v}_i corresponding to its children or-nodes U_i . The response of the and-node A_t is defined as,

$$R_t^a(\Lambda_t, V_t) = \omega^a \cdot \phi^a(\Lambda_t) + \sum_{i \in ch(t)} R_i^u(\Lambda_t, \mathbf{v}_i), \qquad (3)$$

where $\phi^a(\Lambda_t)$ is the BOW histogram globally extracted from the 3-D volume Λ_t centered at I_t . The second term aggregates the response scores from all or-nodes of A_t 's children.

Root-node: The root-node is a global potential function that verifies the temporal compatibility of model, including three terms: the global BoW histogram of the video clip, aggregated scores of its children and-nodes, and temporal displacements of anchor frames. Fig. 4 illustrates the temporal composition by the root-node. We employ the root-node for searching for the best localizations of T anchor frames. We introduce the latent variable Δ_t to indicate the temporal displacement of each anchor frame I_t , which will be calculated during inference. This implicitly carries the temporal ordering constraints which are crucial for discriminating human activities.

In particular, the temporal displacement penalty ξ_t punishes the position of the and-node A_t (*i.e.* one anchor frame) shifting far away from the initial anchor point τ_t in time. Once Δ_t is optimized, the position of each anchor frame can be determined by $\tau_t + \Delta_t$ accordingly. We define ξ_t by,

$$\xi_t = -\omega_t^\tau \cdot \Delta_t, \tag{4}$$

where ω_t^{τ} is the corresponding parameter. The response of the rootnode can be then defined as,

$$R^{r}(X, \mathbf{V}, \mathbf{\Delta}) = \omega^{r} \cdot \phi^{r}(X) + \sum_{t=1}^{r} R^{a}_{t}(\Lambda_{t}, V_{t}) + \xi_{t}, \quad (5)$$

where $\phi^r(X)$ is the BoW histogram feature extracted from the whole video sequence X. $\mathbf{V} = (V_1, \dots, V_T)$ and $\mathbf{\Delta} = (\Delta_1, \dots, \Delta_T)$ are latent variables in the model for specifying the spatial and temporal configurations.

3.2 Contextual Interactions

Spatial interactions. We impose spatial contextual interactions, *i.e.* spatial edges, between pairwise spatially adjacent leaf-nodes

in each anchor frame, as Fig. 2(c) illustrates. Note that we only link the edges between a pair of leaf-nodes that are respectively associated with two different or-nodes.

For one edge connecting two leaf-nodes $(L_j, L_{j'})$, we define it with a set of informative relations, *i.e.* a 8-bin binary feature $\varphi^s(L_j, L_{j'})$: above, below, left, right, near, far, clockwise and anti-clockwise between two adjacent leaf-nodes. The relations are visualized in Fig. 3. Suppose one edge connects two leaf-nodes $(L_j, L_{j'})$ which detect action parts at positions p_j and $p_{j'}$ respectively. The centered red rectangle represents the location p_j , and the other red rectangles represent the adjacent parts. In the right chart of Fig. 3, the dashed line represents the initial layout of the two leaf-nodes, and the black solid line the adjusted actual layout during inference. Then we define the relations as,

- near or far: If p_{j'} is fallen into the outer dashed ellipse, it is near to p_j, *i.e.* the bin of near is activated (*i.e.* set as 1); otherwise it is far to p_j.
- *above, below, left* or *right*: The corresponding bin is set as 1 only if the center of $p_{j'}$ is inside the corresponding dashed rectangles.
- *clockwise* or *anti-clockwise*: one of the two relations is activated (*i.e.* set as 1) according to the angle between the dashed line and the black solid line.

The relations intuitively encode the spatial contexts of two action parts detected via the two leaf-nodes with respect to two different or-nodes. The response of the pairwise potentials can be parameterized as,

$$\Gamma_{jj'}^s = \beta_{jj'}^s \cdot \varphi^s(L_j, L_{j'}), \tag{6}$$

where $\beta_{jj'}^{s}$ is the corresponding 8-bin parameter vector.

Temporal interactions. We also impose the edges in temporal domain in our model to represent the temporal interactions of action parts. The edges connect temporally adjacent leaf-nodes, illustrated in Fig. 4. The edges are connected between any pair of leaf-nodes $(L_j, L_{j'})$ that belong to the same part within the two adjacent anchor frames respectively. A set of temporal relations are collected to concatenate a 4-bin binary feature vector $\varphi^{\tau}(L_j, L_{j'})$.

Specifically, we adopt four predicates: *intersect, after, meets, interrupt*, inspired by Allen's temporal predicates [2] [26]. These predicates describe relations between two time intervals. The action part detected by one leaf-node L_j for a specific anchor frame is described by the feature ϕ^l extracted from a 3-D volume with time span ρ . Note that we ignore some predicates of ordering such as *before* and *equals*, as the order of temporally adjacent anchor frames is supposed to be fixed. Assume that leaf-node L_j is associated with the anchor frame localized at $\tau_t + \Delta_t$, (initial position plus displacement). The starting and ending time $(f_{L_j}^{start}, f_{L_j}^{end})$ of L_j can be calculated as,

$$f_{L_j}^{start} = \tau_t + \Delta_t - \frac{\rho}{2},$$

$$f_{L_j}^{end} = \tau_t + \Delta_t + \frac{\rho}{2}.$$
(7)

Then we can define the four temporal predicates for the two temporal adjacent leaf-nodes $(L_j, L_{j'})$ as,

$$intersect(L_{j}, L_{j'}) \iff f_{L_{j'}}^{start} < f_{L_{j}}^{end},$$

$$after(L_{j}, L_{j'}) \iff f_{L_{j}}^{end} < f_{L_{j'}}^{start} < f_{L_{j}}^{end} + \rho,$$

$$meets(L_{j}, L_{j'}) \iff f_{L_{j'}}^{start} = f_{L_{j}}^{end},$$

$$interrupt(L_{j}, L_{j'}) \iff f_{L_{j}}^{end} + \rho < f_{L_{j'}}^{start}.$$
(8)



Figure 4: Illustration of temporal compositions. The input video is decomposed into a number of discrete temporal anchor frames. The optimal position of each anchor frame is localized in $\Delta_t + \tau_t$, *i.e.* the temporal displacement Δ_t plus the initial anchor point τ_t . The temporal contextual interactions are defined in temporally adjacent action parts.

Thus, we define the response of one temporal edge linking two leafnodes accordingly,

$$\Gamma_{jj'}^{\tau} = \beta_{jj'}^{\tau} \cdot \varphi^{\tau}(L_j, L_{j'}), \qquad (9)$$

where $\beta_{jj'}^{\tau}$ is the corresponding 4-bin parameter. If the pairwise leaf-nodes $(L_j, L_{j'})$ satisfies the specific predicate, the corresponding bin is set to 1, otherwise 0.

Therefore, the overall response of the STAOG model is:

$$R^{g}(X, \mathbf{V}, \mathbf{\Delta}) = R^{r}(X, \mathbf{V}, \mathbf{\Delta}) + \sum_{j}^{n} (\sum_{j' \in \gamma^{s}(j)} \Gamma_{jj'}^{s} \cdot v_{j} \cdot v_{j'} + \sum_{j' \in \gamma^{\tau}(j)} \Gamma_{jj'}^{\tau} \cdot v_{j} \cdot v_{j'}),$$
(10)

where $(\mathbf{V}, \boldsymbol{\Delta})$ are the hidden variables in the STAOG model. The second term defines the spatio-temporal contextual interactions between the leaf-nodes. $\gamma^s(j)$ is denoted as the set of leaf-node L_j 's neighbors which are spatially adjacent to the leaf-node L_j , and $\gamma^{\tau}(j)$ is introduced for the leaf-nodes which are temporally adjacent to the leaf-node L_j . Intuitively, spatial interactions between action parts guarantee the spatial coherence, as well as temporal interactions embedding the temporal contextual relations. Briefly, we refer $\mathcal{L} = (\mathbf{V}, \boldsymbol{\Delta})$ as the latent variables in the following. The Eq.10 can be briefly written as:

$$R^{g}(X,\mathcal{L}) = \psi \cdot \Phi(X,\mathcal{L}), \qquad (11)$$

where ψ includes the complete parameters of the STAOG model, and $\Phi(X, \mathcal{L})$ denotes the overall feature vector.

4. INFERENCE

The inference task is to detect T optimal temporal anchor frames for one video instance as well as the spatial composition of action parts within each anchor frame. In our approach, we perform a cascaded search that integrates three steps: spatial testing, temporal testing and global verification to maximize the global potential $R^g(X, \mathcal{L})$ defined in Eq.10.

Step 1. Spatial Testing via the and-nodes.

The subgraph of the STAOG model, rooted at one and-node, can be viewed as the spatial composition classifier for localizing action parts in one frame. We first use all existing leaf-nodes to search for candidate actions parts. Assume the leaf-node L_j associated with the frame I_t detects the action part at the position p_j^* by maximizing the response in Eq.1. Each or-node is allowed to activate only one leaf-node, then a possible configuration consisting of action parts is decided by the indicator variables of the or-node, (*i.e.* \mathbf{v}_i for ornode U_i , indicating which leaf-node is activated). In this way, a set of possible configuration hypotheses $\{V_t\}$ are generated for further testing, which ensemble the hypotheses proposed by the or-nodes for the frame I_t . In practice, we limit the maximum number of hypotheses by setting a threshold on $R_t^a(\Lambda_t, V_t)$ in Eq.3.

Algorithm 1 Inference Algorithm

Input:

A learned STAOG model G, the action parts detected by all leafnodes by maximizing the response in Eq.1.

Initialization:

The set of possible hypotheses $l_t = \{\}$ for all $t \in 1, ..., T$ anchor frames.

Iteration:

for all $t = 1 \cdots T$ do For each and-node A_t , a set of temporal displacement steps Σ is predefined for sliding the possible anchor frames.

(a) initialize the set of pair terms
$$Q = \{Q_1, \dots, Q_K\}$$

for all or-nodes of A_t 's children.

(b) generate a set of pair terms
$$Q_i$$
 for each or-node U_i .

for all
$$U_i, i \in ch(t)$$
 do
for all $L_{j_2} j \in ch(i)$ do

$$Q_i = Q_i \cup (i, j)$$
.
end for

end for

- (c) obtain possible hypotheses $V_t = (\mathbf{v}_1, \dots, \mathbf{v}_K)$ by assembling the indicator variables of K or-nodes according to the set Q.
- (d) The set of possible hypotheses for each specific displacement Δ_t is constructed as $l_t = l_t \cup (\{V_t\}, \Delta_t)$. end for

end for

Assemble these hypothesis $\{l_t\}$ for all T anchor frames orderly to generate the set of hypotheses sequence l. Each possible configuration $(\mathbf{V}, \boldsymbol{\Delta})$ belongs to the set l. The global response of STAOG model can be calculated by Eq.12. **Output:**

The latent variables $\mathbf{V}, \boldsymbol{\Delta}$ and the final score $S_{\psi}(X)$.

Step 2. Temporal Testing via the root-node.

We apply the spatial testing with the and-nodes to localize a number of candidate anchor frames over several frames. The scores over candidate anchor frames (proposed by the and-nodes) are aggregated via the root-node for a possible temporal composition. For efficiency, we utilize a fixed number of discrete steps Σ for searching each anchor frame. Several possible hypotheses are then produced with different anchor frame determinations by sliding the discrete steps Σ . In addition, we re-weight the hypotheses at the root-node in Eq.5 by considering the temporal displacements of anchor frames as well as the the global features over the video clip. Intuitively, the hypotheses are represented by the specified latent variables ($\mathbf{V}, \boldsymbol{\Delta}$).

Step 3. Global Verification.

Given all the hypotheses from the root-node, we apply the global potential function defined in Eq.10 to validate the optimal detection. The objective of the global verification is to cope with the noisy local detections on leaf-nodes. It combines the score of the

(10)

root-node with the responses of spatial and temporal contextual interactions (edges).

The optimal response $S_{\psi}(X)$ of the model as well as the latent variables $(\mathbf{V}, \boldsymbol{\Delta})$ can be calculated as,

$$S_{\psi}(X) = \max_{\mathbf{V}, \mathbf{\Delta}} (\psi \cdot \Phi(X, \mathbf{V}, \mathbf{\Delta})).$$
(12)

Algorithm 1 summarizes the overall algorithm of the inference.

5. STRUCTURAL LEARNING

We formulate the structural learning of STAOG model as a joint optimization task for model structure and parameters. We solve this model by a novel latent learning method extended from the CCCP framework [43]. This algorithm iterates to train the model in a dynamic manner: the leaf-nodes can be automatically created or removed to reconfigure the model structure. The model structure is determined by latent variables $\mathcal{L} = (\mathbf{V}, \boldsymbol{\Delta})$ that are inferred in each step.

Let $D = ((X_1, y_1), (X_2, y_2), \dots, (X_N, y_N))$ be a set of labeled training samples, with $y_k \in \{1, -1\}$. The feature vector for each sample (X, y) is defined as,

$$\Phi(X, y, \mathcal{L}) = \begin{cases} \Phi(X, \mathcal{L}) & \text{if } y = +1, \\ 0 & \text{if } y = -1. \end{cases}$$
(13)

The temporal anchor frames and spatial configurations for each frame can be optimized by maximizing $R^g(X, \mathcal{L})$ in the inference procedure. We refer to $\mathcal{L} = (\mathbf{V}, \boldsymbol{\Delta})$ as the latent variables, then we redefine Eq.12 as,

$$S_{\psi}(X) = \max_{\substack{y,\mathcal{L}}} (\psi \cdot \Phi(X, y, \mathcal{L})), \tag{14}$$

The optimization of this function can be solved by the latent structural SVM framework,

$$\min_{\psi} \frac{1}{2} \|\psi\|^2 + C \sum_{k=1}^{N} [\max_{y,\mathcal{L}} (\psi \cdot \Phi(X_k, y, \mathcal{L}) + h(y_k, y)) - \max_{x} (\psi \cdot \Phi(X_k, y_k, \mathcal{L}))],$$
(15)

where C is a penalty parameter set as 0.003 empirically and $h(y_k, y)$ is the cost function, where $h(y_k, y) = 0$ if $y_k = y$, otherwise 1. The optimization problem described above is not convex in general. Following the CCCP framework, we convert the function in Eq.15 into a convex and concave form as,

$$\begin{split} \min_{\psi} \left[\frac{1}{2} \|\psi\|^{2} + C \sum_{k=1}^{N} \max_{y,\mathcal{L}} (\psi \cdot \Phi(X_{k}, y, \mathcal{L}) + h(y_{k}, y)) \right] \\ - \left[C \sum_{k=1}^{N} \max_{\mathcal{L}} (\psi \cdot \Phi(X_{k}, y_{k}, \mathcal{L})) \right] \\ = \min_{\psi} [f(\psi) - g(\psi)], \end{split}$$
(16)

where $f(\psi)$ represents the first two terms, and $q(\psi)$ the last term. This leads to an iterative learning algorithm that alternates estimating model parameters and the hidden variables \mathcal{L} . However, we still need to dynamically determine the graph configuration, *i.e.* the production of leaf-nodes associated with or-nodes. An additional step for dynamically reconfiguring structure is added between the two original steps. The procedure is presented as follows.

(I) The model parameter ψ_t obtained in the previous iteration is fixed. We find a hyperplane q_t to upper bound the concave part $g(\psi)$ in Eq. 16. Specifically, q_t is the derivative of $g(\psi)$. Thus, we have $-g(\psi) \leq -g(\psi_t) + (\psi - \psi_t) \cdot q_t, \forall \psi$. The optimal latent

Algorithm 2 Learning algorithm for STAOG model

Input:

2.

3.

Training samples, $D = ((X_1, y_1), (X_2, y_2), \dots, (X_N, y_N)).$ Output:

The trained STAOG model.

- Initialization:
 - 1. Initialize the positions of action parts and anchor frames for all samples.
- 2. Initialize the latent variables \mathcal{L} and parameters ψ . repeat

 - Estimate the latent variable $\mathcal{L}_k^* = (\mathbf{V}_k^*, \mathbf{\Delta}_k^*)$ for each pos-itive sample (X_k, y_k) during inference. 1.
 - (a) Localize the anchor frames and action parts using the current latent variables $(\mathbf{V}^*, \boldsymbol{\Delta}^*)$.
 - (b) For each or-node U_i , we obtain a set of feature vectors of its children leaf-nodes for positive samples, and regroup the feature vectors by Spectral clustering.
 - (c) Reconfigure leaf-nodes according to clustering results to generate a new structure. Calculate the energy Eq. 17 as $E(\psi_t^d)$.

- if
$$E(\psi_t^d) < E(\psi_t)$$
 then

Accept the new model structure and estimate the parameters $\psi_{t+1} = argmin_{\psi}(f(\psi + \psi \cdot q_t^d)).$

keep the previous model structure. Estimate the parameters
$$\psi_{t+1} = argmin_{\psi}(f(\psi + \psi \cdot q_t))$$
.

– end if

until The optimization function defined in Eq. 15 converges.

variable \mathcal{L}_k^* is calculated by $\mathcal{L}_k^* = argmax_{\mathcal{L}}(\psi_t \cdot \Phi(X_k, y_k, \mathcal{L}))$ for each positive example. Then the hyperplane is constructed as $q_t = -C \sum_{k=1}^{N} \Phi(X_k, y_k, \mathcal{L}_k^*).$

(II) In the second step, the STAOG model is adjusted by structural reconfiguration with applying the current model on training examples. The reconfiguration is performed for each part, i.e. ornode, independently, with the fixed latent variable $\mathcal{L}_{k}^{*} = (V_{k}^{*}, \Delta_{k}^{*})$. Note that each action part detected by one leaf-node is mapped to several feature bins at specific positions in the vector $\Phi(X_k, y_k, \mathcal{L}_k^*)$.

For each or-node U_i , we apply its children leaf-nodes for detecting parts in all positive samples. Assume that leaf-node L_i detects an action part on the k-th sample using the feature vector π_i^k , which is a sub-vector of the complete feature vector of k-th sample, $\Phi(X_k, y_k, \mathcal{L}_k^*)$. And we obtain a set of feature vector $\{\pi_i^k\}$ for all samples. The vectors detected by the same leaf-node are first grouped into one cluster, *i.e.* one cluster for one leaf-node. We denote the cluster for the *j*-th leaf-node as Ω_j . Then we perform the spectral clustering algorithm with the Euclidean distance on vectors of all leaf-nodes of or-node U_i 's children for all positive samples, and the similar vectors are grouped together. We re-arrange the feature vectors for all samples based on the newly generated partition. For example, if the feature vector π_j^k is grouped from Ω_j into another cluster Ω'_j , we adjust the position of π^k_j in $\Phi(X_k, y_k, \mathcal{L}^*_k)$, *i.e.* by moving the feature bins into the position representing j'th leaf-node. If Ω'_i is a newly generated cluster, we thus create a new leaf-node accordingly. By analogy, we remove one leaf-node if few samples are grouped into the corresponding cluster. In this way, the structure of U_i is reconfigured with the feature vector rearrangement. In practice, we constrain the extent of structural reconfiguration, *i.e.* only few leaf-nodes can be created or removed in one iteration. We present a toy example in Fig.5 for illustration. In Fig.5(a), a leaf-node associated with the or-node U_7 is created to better handle the intra-class variance; A leaf-node is removed if there is another similar one, (e.g. the leaf-node associated with the



Figure 5: Illustration of discriminative structural learning. We reconfigure the model structure by re-arranging the feature vector, as the example illustrates. Parts of the STAOG model reconfigured in two iterations are shown in (a), where the left one represents the original model and the other one the new model. During this step, a new leaf-node associated with U_7 and U_{11} is created and a leaf-node associated with U_8 is removed. Assume that we use 5 samples, X_1, \ldots, X_5 , for the structure learning. (b) shows the feature vectors detected by the same leaf-node are first grouped into one cluster, *i.e.* one cluster for one leaf-node. (c) illustrates the feature rearrangement after clustering. For example, the feature vector of sample X_1 is grouped from cluster Ω_9 into cluster Ω_{13} , we move the feature bins π_9^1 into the bins corresponding to the leaf-node L_9 . Cluster Ω_{14} is a newly generated cluster, we thus create a new leaf-node accordingly.

or-node U_8). The sub-vector of π_9^1 of sample X_1 is grouped from cluster Ω_9 to cluster Ω_{13} , then the feature bins are moved from π_9^1 to π_{13}^1 as Fig.5(b) shows.

After this structure reconfiguration, we obtain the new feature vector for each sample, $\Phi^d(X_k, y_k, \mathcal{L}_k^*)$, and the hyperplane is recalculated as $q_t^d = -C \sum_{k=1}^N \Phi^d(X_k, y_k, \mathcal{L}_k^*)$, accordingly.

(III) The newly generated model structure can be represented by the feature vector $\Phi^d(X_k, y_k, \mathcal{L}_k^*)$. The model parameters can be learned by solving $\psi_t^d = argmin_{\psi}(f(\psi + \psi \cdot q_t^d))$. The optimization task in Eq.15 becomes,

$$\min_{\psi} \frac{1}{2} \|\psi\|^2 + C \sum_{k=1}^{N} [\max_{y,\mathcal{L}} (\psi \cdot \Phi(X_k, y, \mathcal{L}) + h(y_k, y)) - \psi \cdot \Phi^d(X_k, y_k, \mathcal{L}_k^*))].$$
(17)

This is a standard structural SVM problem, which can be solved in the cutting plane method and Sequential Minimal Optimization. The energy Eq.17 can be calculated by $E(\psi_t^d) = f(\psi_t^d) - g(\psi_t^d)$. We accept the new model structure until $E(\psi_t^d) < E(\psi_t)$ and $\psi_{t+1} = \psi_t^d$. Otherwise, we keep the model structure as in the previous iteration and the parameter vector is calculated by $\psi_{t+1} = argmin_{\psi}(f(\psi + \psi \cdot q_t))$. Thus, we ensure that the optimization function will decrease in each iteration. We repeat the 3-step iteration until convergence. Algorithm 2 summarizes the overall algorithm of learning a STAOG model. In the case of multi-class classification, we use a one-against-rest approach and select the class with the highest score.

6. EXPERIMENTS

We test our STAOG model on two different action recognition databases: UCF YouTube [21] and Olympics Sports [24]. The video resolution is normalized to 320×240 . The YouTube dataset contains 11 action categories, which is challenging due to large variation in camera motion, object appearance, pose/view and large intra-class variability. We follow the standard setup using leave-one-out cross validation for a pre-defined set of 25 folds.Average

accuracy over all classes is reported as performance measure. The Olympic Sports dataset consists of 16 different sports classes that contain complex motions going beyond simple punctual or repetitive actions. The challenges of Olympic sports arise from background clutters, viewpoints and complex sequence of primitive actions. Each action is performed only by a single actor and represents a temporal sequence of primitive actions (*e.g.* triple-jumping, pole-vault and diving). We use the same train-test split setup and the average precision (AP) for each of the action classes as in [35].

6.1 Implementation

We fix the number of and-nodes (anchor frames) in the STAOG model as T = 3 for UCF YouTube dataset, and T = 5 for Olympic Sports dataset empirically. We need more anchor frames for Olympic dataset because the actions are more complex and last over more frames. The parameter T can be roughly estimated by the action temporal complexity in general. The number of spatial layout for each anchor frame is fixed as 2×2 , and thus K = 4 or-nodes for each anchor frame. There are at most m = 4 leaf-nodes associated to each or-node. We extract the interest points described by the HOG and HOF features by utilizing the code published in [17] beforehand. The size of one action part (within a 3-D volume) is empirically set to 60×60 pixels spanning $\rho = 15$ frames. The dimension of the generated dictionary is set as 300 for describing action parts in each anchor frame and the global features of the anchor frames. We set the discrete temporal steps searching for anchor frames: $\Sigma = [\pm 2, \pm 4, \pm 6, \pm 8, \pm 10]$. The convergence of our learning algorithm usually takes $9 \sim 10$ iterations.

The experiments are carried out on a PC with Core I5 3.0GHZ CPU and 4GB memory. The average CPU-time used to process a video from Olympic Sports dataset is 200 seconds and 150 for a video in the UCF YouTube dataset. In particular, it takes 7 seconds for processing one frame in Olympic Sports dataset, and 4 seconds for each frame in the UCF YouTube dataset on average. The efficiency of our method is slightly diverse with the densities of the feature points in the video sequence.

	Liu et al [21].	Zhang et al [44].	Ikizler-Cinbis et al [12].	Dense trajectories [36].	Ours-1	Ours-2	Ours(full)
b_shoot	53.0%	98.0%	48.5%	43.0%	58.4%	62.0%	77.9%
bike	73.0%	74.0%	75.2%	91.7%	82.1%	87.3%	88.6%
dive	81.0%	80.0%	95.0%	99.0%	98.4%	98.6%	98.8%
golf	86.0%	68.0%	95.0%	97.0%	95.7%	95.3%	97.4%
h_ride	72.0%	65.0%	73.0%	85.0%	81.3%	86.0%	88.0%
s_juggle	54.0%	67.0%	53.0%	76.0%	66.0%	81.6%	82.2%
swing	57.0%	71.0%	66.0%	88.0%	85.2%	84.1%	85.4%
t_swing	80.0%	68.0%	77.0%	71.0%	69.6%	80.0%	80.7%
t_jump	79.0%	80.0%	93.0%	94.0%	90.2%	94.7%	95.8%
v_spike	73.3%	77.0%	85.0%	95.0%	89.7%	90.6%	96.4%
walk	75.0%	54.0%	66.7%	87.0%	85.2%	86.2%	87.4 %
Accuracy	71.2%	72.9%	75.2 %	84.2%	82.0%	86.0%	88.9%

Table 1: Accuracy per action class and average accuracy for all classes on the YouTube dataset [21].



Figure 6: Example inference results on three different action models *e.g.* basketball-layup (a), clean-and-jerk (b) and triple-jump(c) learned on the Olympic Sports dataset and each of action category includes two instances. The red boxes in each frame represent the discovered discriminative action parts. Our model successfully localizes the accurate anchor frames across the instances in the long action videos. In addition, it is noticed that the large intra-class variabilities can be captured by our model.

	Niebles et al [24].	Tang et al [35].	Ours(full)
h-jump	27.0%	18.4%	35.6%
1-jump	71.7%	81.8%	86.4%
t-jump	10.1%	16.1%	36.2%
p-vault	90.8 %	84.9%	84.3%
g-vault	86.1%	85.7%	83.1%
s-put	37.3%	43.3%	56.8%
snatch	54.2%	88.6%	89.0%
c-jerk	70.6%	78.2%	83.3%
j-throw	85.0%	79.5%	78.1%
h-throw	71.2%	70.5%	75.4%
d-throw	47.3%	48.9%	53.3%
d-platform	95.4 %	93.7%	92.8%
d-board	84.3%	79.3%	76.5%
basketball	82.1%	85.5%	86.7%
bowling	53.0%	64.3%	62.0%
t-serve	33.4%	49.6%	62.3%
mAP	62.5%	66.8%	71.4%

 Table 2: Average Precision(AP) values on the Olympic Sports dataset [24].

6.2 **Results and Comparisons**

Compared with the recently proposed methods on the YouTube dataset, our model outperforms the state-of-the-art: we achieve the accuracy of 88.9% in YouTube dataset, the reported results of the competing algorithms are: 71.2% in [21], 72.9% in [44],75.2% in [12] and 84.2% in [36]. The accuracy scores for all categories are reported in Table 1. Our method outperforms on 7 out of the 11 categories which have relatively large intra-class variance or background disturbance. In the Olympic Sports dataset, we obtain better AP scores for 10 out of the 16 categories, and overall AP score 71.4%, better than the previous methods [35, 24]. The competing method proposed by Tang et al. [35] utilizes the variable-duration HMM for learning the temporal structure in the video. Our results show that our compositional model with the explicit spatial and temporal relations can achieve better performance. The detailed results are reported in Table 2.

Figure. 6 illustrates the inference results on three action categories from the Olympic Sports dataset, each of which includes two instances. Our model also localizes the action parts in the anchor frames, as they are discriminative in appearance and motion. These results demonstrate well the capability of our model, since the scenarios of actions contain the very realistic challenges in video action recognition. The spatial compositions defined over the and-nodes enable us to handle pose/view variations and background disturbances. The temporal compositions are effective to localize anchor frames in videos against various motion frequency, temporal locations and video length.

For further evaluation, we conduct three empirical analysis in different model settings as follows.

(I) We simplify the temporal compositions by discarding the displacement Δ_t for each anchor frame, *i.e.* fixing the temporal structure. We report AP scores of this model setting in the fourth column of Table 1, named as "Ours-1". The average accuracy is 82%, 6.9% less than our complete model.

(II) We also evaluate the benefit of spatio-temporal contextual interactions. Our model can be simplified into a tree structure by removing the interactions. The accuracies are shown in the sixth column of Table 1, named as "Ours-2". We can observe that the spatio-temporal contextual interactions make the accuracies increase 2.9% on average. The increased performance of the interac-



Figure 7: Empirical analysis for different settings of spatial compositions, where we set different maximum numbers m of leaf-nodes under the or-nodes. Each pillar represents the accuracy for one action category. The color indicates the results with different settings, m = 2 or m = 4.

tions also speaks in favour for our model, as it shows that through better associations between the anchor frames and action parts, it is possible to achieve even better accuracies.

(III) One may be interested in how the performances are improved by introducing the or-nodes in spatial compositions, which is one of the key components in the STAOG model. In this experiment, we set different maximum numbers m of leaf-nodes under the or-nodes, *i.e.* how many leaf-nodes at most can be created for the model. We compare the results on the YouTube dataset in Fig. 7, and observe that m = 4 achieves the better results in general than m = 2, 82.4%. In practice, the number m is not sensitive, as the exact number of leaf-nodes is decided by the clustering on data during the structural learning.

7. CONCLUSION

This paper studies a novel hierarchical model for human action recognition, in the form of a configurable Spatio-Temporal And-Or Graph. This model is shown to handle well realistic challenges in action recognition. Moreover, we consider two aspects to improve our method. First, the model can be integrated with high level semantic information to represent multi-agent complex events. Second, we plan to speed up the algorithm for large-scale processing.

8. REFERENCES

- [1] M. Albanese, R. Chellappa, N. P. Cuntoor, V. Moscato, A. Picariello, V. S. Subrahmanian, and O. Udrea. A constrained probabilistic petri net framework for human activity detection in video. *IEEE Transactions on Multimedia*, 10(6):982–996, 2008.
- [2] J. F. Allen and G. Ferguson. Actions and events in interval temporal logic. J. Log. Comput., 4(5):531–579, 1994.
- [3] M. R. Amer, D. Xie, M. Zhao, S. Todorovic, and S. C. Zhu. Cost-sensitive top-down/bottom-up inference for multiscale activity recognition. In *ECCV* (4), pages 187–200, 2012.
- [4] W. Brendel and S. Todorovic. Learning spatiotemporal graphs of human activities. In *ICCV*, pages 778–785, 2011.
- [5] L. Cao, Z. Liu, and T. S. Huang. Cross-dataset action detection. In *CVPR*, pages 1998–2005, 2010.
- [6] W. Choi, K. Shahid, and S. Savarese. Learning context for collective activity recognition. In *CVPR*, pages 3273–3280, 2011.
- [7] P. Dollár, V. Rabaud, G. Cottrell, and S. Belongie. Behavior recognition via sparse spatio-temporal features. In VS-PETS,

October 2005.

- [8] X. Duan, L. Lin, and H. Chao. Discovering video shot categories by unsupervised stochastic graph partition. *IEEE Transactions on Multimedia*, 15(1):167–180, 2013.
- [9] T. V. Duong, H. H. Bui, D. Q. Phung, and S. Venkatesh. Activity recognition and abnormality detection with the switching hidden semi-markov model. In *CVPR (1)*, pages 838–845, 2005.
- [10] P. F. Felzenszwalb, R. B. Girshick, D. A. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(9):1627–1645, 2010.
- [11] A. Gupta, P. Srinivasan, J. Shi, and L. S. Davis. Understanding videos, constructing plots learning a visually grounded storyline model from annotated videos. In *CVPR*, pages 2012–2019, 2009.
- [12] N. Ikizler-Cinbis and S. Sclaroff. Object, scene and actions: Combining multiple features for human action recognition. In ECCV (1), pages 494–507, 2010.
- [13] L. Jiang, A. G. Hauptmann, and G. Xiang. Leveraging high-level and low-level features for multimedia event detection. In ACM Multimedia, pages 449–458, 2012.
- [14] Y.-G. Jiang, Q. Dai, X. Xue, W. Liu, and C.-W. Ngo. Trajectory-based modeling of human actions with motion reference points. In *ECCV* (5), 2012.
- [15] A. Kläser, M. Marszalek, and C. Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC*, 2008.
- [16] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64(2-3):107–123, 2005.
- [17] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008.
- [18] Z. Li, Y. Fu, T. S. Huang, and S. Yan. Real-time human action recognition by luminance field trajectory analysis. In *ACM Multimedia*, pages 671–676, 2008.
- [19] L. Lin, X. Wang, W. Yang, and J. Lai. Learning contour-fragment-based shape model with and-or tree representation. In *CVPR*, pages 135–142, 2012.
- [20] 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.
- [21] J. Liu, J. Luo, and M. Shah. Recognizing realistic actions from videos. In CVPR, pages 1996–2003, 2009.
- [22] Z. Ma, Y. Yang, Y. Cai, N. Sebe, and A. G. Hauptmann. Knowledge adaptation for ad hoc multimedia event detection with few exemplars. In *ACM Multimedia*, pages 469–478, 2012.
- [23] M. Marszalek, I. Laptev, and C. Schmid. Actions in context. In CVPR, 2009.
- [24] J. C. Niebles, C.-W. Chen, and F.-F. Li. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV* (2), pages 392–405, 2010.
- [25] M. Raptis, I. Kokkinos, and S. Soatto. Discovering discriminative action parts from mid-level video representations. In *CVPR*, pages 1242–1249, 2012.
- [26] M. S. Ryoo and J. K. Aggarwal. Spatio-temporal relationship match: Video structure comparison for recognition of complex human activities. In *ICCV*, pages 1593–1600, 2009.
- [27] S. Sadanand and J. J. Corso. Action bank: A high-level representation of activity in video. In *CVPR*, pages 1234–1241, 2012.

- [28] S. Satkin and M. Hebert. Modeling the temporal extent of actions. In *ECCV*(1), pages 536–548, 2010.
- [29] 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.
- [30] Z. Si, M. Pei, B. Yao, and S.-C. Zhu. Unsupervised learning of event and-or grammar and semantics from video. In *ICCV*, pages 41–48, 2011.
- [31] C. Snoek, M. Worring, and A. W. M. Smeulders. Early versus late fusion in semantic video analysis. In ACM Multimedia, pages 399–402, 2005.
- [32] Y. Song, L.-P. Morency, and R. Davis. Multi-view latent variable discriminative models for action recognition. In *CVPR*, pages 2120–2127, 2012.
- [33] Y. Song, S. Tang, Y.-T. Zheng, T.-S. Chua, Y. Zhang, and S. Lin. Exploring probabilistic localized video representation for human action recognition. *Multimedia Tools Appl.*, 58(3):663–685, 2012.
- [34] J. Sun, X. Wu, S. Yan, L. F. Cheong, T.-S. Chua, and J. Li. Hierarchical spatio-temporal context modeling for action recognition. In *CVPR*, pages 2004–2011, 2009.
- [35] K. Tang, F.-F. Li, and D. Koller. Learning latent temporal structure for complex event detection. In *CVPR*, pages 1250–1257, 2012.
- [36] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *CVPR*, pages 3169–3176, 2011.
- [37] S. Wang, Y. Yang, Z. Ma, X. Li, C. Pang, and A. G. Hauptmann. Action recognition by exploring data distribution and feature correlation. In *CVPR*, pages 1370–1377, 2012.
- [38] X. Wang and L. Lin. Dynamical and-or graph learning for object shape modeling and detection. In *NIPS*, pages 242–250, 2012.
- [39] Y. Wang and G. Mori. Hidden part models for human action recognition: Probabilistic versus max margin. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(7):1310–1323, 2011.
- [40] B. Yao and S. C. Zhu. Learning deformable action templates from cluttered videos. In *ICCV*, pages 1507–1514, 2009.
- [41] G. Yu, J. Yuan, and Z. Liu. Real-time human action search using random forest based hough voting. In ACM Multimedia, pages 1149–1152, 2011.
- [42] J. Yuan, Z. Liu, and Y. Wu. Discriminative video pattern search for efficient action detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(9):1728–1743, 2011.
- [43] A. L. Yuille and A. Rangarajan. The concave-convex procedure. *Neural Computation*, 15(4):915–936, 2003.
- [44] Y. Zhang, X. Liu, M.-C. Chang, W. Ge, and T. Chen. Spatio-temporal phrases for activity recognition. In *ECCV* (3), pages 707–721, 2012.
- [45] X. Zhou, X. Zhuang, S. Yan, S.-F. Chang, M. Hasegawa-Johnson, and T. S. Huang. Sift-bag kernel for video event analysis. In ACM Multimedia, pages 229–238, 2008.
- [46] G. Zhu, M. Yang, K. Yu, W. Xu, and Y. Gong. Detecting video events based on action recognition in complex scenes using spatio-temporal descriptor. In *ACM Multimedia*, pages 165–174, 2009.
- [47] S. C. Zhu and D. Mumford. A stochastic grammar of images. Foundations and Trends in Computer Graphics and Vision, 2(4):259–362, 2006.