Recognition of Group Activities in Videos Based on Single- and Two-Person Descriptors

Stéphane Lathuilière, Georgios Evangelidis, Radu Horaud

► To cite this version:

Stéphane Lathuilière, Georgios Evangelidis, Radu Horaud. Recognition of Group Activities in Videos Based on Single- and Two-Person Descriptors. IEEE Winter Conference on Applications of Computer Vision, Mar 2017, Santa Rosa, CA, United States. pp.217-225, 10.1109/WACV.2017.31 . hal-01430732

HAL Id: hal-01430732

<https://hal.inria.fr/hal-01430732>

Submitted on 10 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recognition of Group Activities in Videos Based on Single- and Two-Person Descriptors *

Stéphane Lathuilière¹, Georgios Evangelidis^{1,2}, Radu Horaud¹

¹INRIA Grenoble, France ²DAQRI Int. Dublin, Ireland

stephane.lathuiliere@inria.fr

Abstract

Group activity recognition from videos is a very challenging problem that has barely been addressed. We propose an activity recognition method using group context. In order to encode both single-person description and two-person interactions, we learn mappings from high-dimensional feature spaces to low-dimensional dictionaries. In particular the proposed two-person descriptor takes into account geometric characteristics of the relative pose and motion between the two persons. Both single-person and two-person representations are then used to define unary and pairwise potentials of an energy function, whose optimization leads to the structured labeling of persons involved in the same activity. An interesting feature of the proposed method is that, unlike the vast majority of existing methods, it is able to recognize multiple distinct group activities occurring simultaneously in a video. The proposed method is evaluated with datasets widely used for group activity recognition, and is compared with several baseline methods.

1. Introduction

Human activity recognition has been an active topic in computer vision over the past years. Although considerable research has been devoted to the recognition of single-person [36, 28, 37, 21] and two-person [31, 32] activities, less attention has been paid to group activities. The latter is however important for a number of applications, e.g. video surveillance, video monitoring, human-robot interaction and video indexing, etc.

In this paper we aim to recognize activities that are performed jointly by groups of people, such as *dancing*, *talking* and so forth. More precisely, we propose a model that

Figure 1: High-dimensional feature vectors are mapped onto a low-dimensional space that encodes both single-person and two-person descriptors. Based on these descriptors, a group-activity label is assigned to each person by solving an energy optimization problem.

can handle videos that contain one or even several **distinct** group activities that occur either simultaneously or in sequence. Such a problem cannot be solved by simply considering multiple instances of single-person activities as long as the collective context carries out important information [7]. As recently shown, the inter-person distance, the motion, and the relative pose of group members seem to be important cues for good recognition performance [35, 20, 4]. However, the way these cues should be combined together is crucial for robust recognition.

In this paper we cast the group-activity recognition problem into a structured-output labelling problem that is solved within an optimization framework. A two-stage mapping is used to obtain low-dimensional features and, in turn, to define the unary/pairwise potentials of the energy terms, e.g. Fig. 1. First, spatiotemporal motion features, e.g. [37] are used in conjunction with bounding boxes which are provided in advance (person detection is not addressed in this paper). Second, we carefully investigate the two-person geometry, motion and relative pose, in order to describe high-level cues such as *persons A and B face each other*, or *stand side-by-side*, or *move in the same direction*. These descriptions are further encoded into pose-activity words in an unsupervised manner. Third, such a multi-stage mapping pro-

*Funding from the European Union FP7 ERC Advanced Grant VHIA (#340113) is greatly acknowledged.

vides low-dimensional yet meaningful interaction descriptors that are eventually used (along with single-person descriptors) to define unary and pairwise potentials of an energy function. The parameters of the latter are learned using structured SVM. It is worth noticing that the estimation of person poses needed to compute two-person descriptors is only performed for model training and not for testing (recognition).

The remainder of the paper is organized as follows. Sec. 2 discusses the related work. In Sec. 3, the energy function is presented. Sec. 4 details the pairwise interaction description, while model learning is explained in Sec. 5. Sec. 6 discusses the features that encode the raw data and Sec. 7 presents the experimental results. Finally, Sec. 8 concludes this work.

2. Related Work

Early human action recognition methods could only handle simple activities performed by a single person in controlled environments, e.g. [34]. More complex scenarios have been subsequently addressed, such as presence of occlusions, changing illumination conditions, moving cameras, dynamic background, e.g. [23]. We note that challenging single-person activities can be dealt with using sophisticated feature representations, e.g. [37]. At the same time, significant effort has been put toward using contextual information to improve activity and object recognition [29, 24]. Such strategies benefit from using the global image content, thus not suffering from low-quality appearance, small objects, or occlusions. The object-action context is addressed in [17, 23, 11, 38, 15] while spatial coherence constraints may be enforced as well [11].

When it comes to group activity, the recognition performance benefits from a cross-person context [2, 27, 30, 19, 35]. As a first attempt, [7] showed that the collective behavior, via modeling relative poses of persons, improves the classification of group activities. While this framework can deal with multiple group activities that are performed simultaneously, the classification is still done individually, i.e. for each detected person. Subsequent work focused on the simpler scenario of a single group activity per frame. Any person interaction description can then be simply integrated in the model, e.g. [19, 35, 12] and not only in feature encoding. This led to the introduction of pairwise potentials within energy-based formulations, so that persons involved in the same activity are jointly considered. Group structure analysis and selection of meaningful pairs of persons [20, 35] seem to further help at the expense of higher computational complexity. However, these methods relies on tracking each person, that can be inaccurate in crowded scenarios. Note that [4] models groups of humans in still

images. More recently, deep structured models have been used for single group activity [10, 14]. While these methods reach satisfying results, their generalization to the multiple activity case is not straightforward. Actually, when several activities can occur simultaneously, the size of the label space grows exponentially with the number of persons. Thus, complex or restrictive formulations cannot be used in practice. Temporal information has been also used in [8, 5, 16] to deal with issues that arise from the social context, e.g. occlusions, gathering etc. In a similar manner, [22] considers temporal interactions between persons to analyse team sports. In some sense, however, these two strategies contradict each other since the finer the group modelization, the more difficult its integration into a temporal framework. Conversely, adding the time dimension into graph formulations leads to computationally demanding solutions, in particular when several groups of different activities appear at the same time. As a consequence, one may need to properly combine these two approaches.

Instead, we propose to investigate the two-person (or person-person) geometry, the motion and the relative poses in order to describe precisely the interactions. The proposed model exploits temporal information only when building descriptors, thus avoiding inference on temporal models. Moreover, the proposed descriptors can be used efficiently and without computationally demanding optimization algorithms in difficult scenarios in which several group activities occurs simultaneously. In addition, and unlike [7, 6, 4], body poses are used for training purposes only, so that the body pose is not required at test time to estimate our pairwise descriptor.

3. Modeling Group Activity

Provided that $I < I_{max}$ persons are detected in an image, we extract high-dimensional single-person feature vectors, $\mathbf{x}_i \in \mathbb{R}^D$, from each bounding box and high-dimensional pairwise feature vectors, $\mathbf{y}_{ij} \in \mathbb{R}^F$, from each pair of boxes, so that the sets $\mathbf{X} = \{\mathbf{x}_i \mid 1 \leq i \leq I\}$ and $\mathbf{Y} = \{\mathbf{y}_{ij} \mid 1 \leq i, j \leq I, i \neq j\}$ are available. The goal is to find a set of activity labels $\mathbf{A} = \{a_i \mid 1 \leq i \leq I\}$ (one label per detected person) where $a_i \in \mathcal{L} = \{l_m \mid 1 \leq m \leq M\}$, that is an M -activity label set. To this end, we define an energy function $E(\mathbf{X}, \mathbf{Y}, \mathbf{A})$ and we seek the optimizer $\mathbf{A}^*$ that best fits with the features:

$$\mathbf{A}^* = \underset{\mathbf{A} \in \mathcal{A}_I}{\operatorname{argmax}} E(\mathbf{X}, \mathbf{Y}, \mathbf{A}), \quad (1)$$

where $\mathcal{A}_I$ is the activity label set. The energy function is defined as:

$$E(\mathbf{X}, \mathbf{Y}, \mathbf{A}) = \sum_{i=1}^I \underbrace{\Psi_1(\mathbf{x}_i, a_i)}_{\text{individual potential}} + \sum_{i \neq j}^{I \times I} \underbrace{\Psi_2(\mathbf{y}_{ij}, a_i, a_j)}_{\text{pairwise potential}} + \underbrace{\Psi_3(\mathbf{A})}_{\text{regularizer}}. \quad (2)$$

The individual potential $\Psi_1(\mathbf{x}_i, a_i)$ models the compatibility of label a_i with $\mathbf{x}_i$; the pairwise potential $\Psi_2(\mathbf{y}_{ij}, a_i, a_j)$ models the compatibility of a pair of labels (a_i, a_j) with $\mathbf{y}_{ij}$; the term $\Psi_3(\mathbf{A})$ enforces grouping, i.e., by assigning the same activity label to multiple persons. A log-linear model leads to the following energy function:

$$E(\mathbf{X}, \mathbf{Y}, \mathbf{A}) = \sum_{i=1}^I \mathbf{w}_{a_i}^{1\top} \phi_1(\mathbf{x}_i) + \sum_{i \neq j}^{I \times I} \mathbf{w}_{a_i, a_j}^{2\top} \phi_2(\mathbf{y}_{ij}) \mathbb{1}(a_i = a_j) + \mathbf{w}^{3\top} \phi_3(\mathbf{A}). \quad (3)$$

The indicator function in the pairwise term makes the sum valid for the pairs of persons that perform the same activity. Grouping is then implicitly considered without adding group variables. Note that the parameters $\mathbf{w}_{a_i}^1, \mathbf{w}_{a_i, a_j}^2$ and $\mathbf{w}^3$ are learned on a training set (Sec. 5). As for the feature functions, they are defined as follows.

Single-person mapping ϕ_1 : Let $\mathbf{B} = \{b_i \mid 1 \leq i \leq I\}$ be the discrete body poses of the detected persons where $b_i \in \Pi = \{\pi_q \mid 1 \leq q \leq Q\} \subset [-\pi, \pi]$, i.e., Π denotes a set of possible discrete body poses. The poses are modeled by an angle between the body orientation and a reference orientation. As in [19], ϕ_1 maps the high-dimensional feature space onto a low-dimensional space of posterior probabilities: $\phi_1(\mathbf{x}_i) = \left(P(b_i = \pi_q, a_i = l_m | \mathbf{x}_i) \right)_{(\pi_q, l_m) \in \Pi \times \mathcal{L}}$. The mapping ϕ_1 is estimated via a linear SVM on training data in the second stage of the pipeline.

Two-person mapping ϕ_2 : This mapping projects a pairwise high-dimensional feature vector onto a low-dimensional space that describes the interaction between two detected persons, e.g., facing each other. The explicit design and estimation of ϕ_2 is one of the main contributions and it is explained in detail in Sec. 4.

Regularization term ϕ_3 : The regularizer is simply defined as $\phi_3(\mathbf{A}) = \left(\mathbb{1}(\nu(\mathbf{A}) = i) \right)_{1 \leq i \leq I_{max}}$, where $\nu(\mathbf{A})$ is the number of different activities in $\mathbf{A}$. For instance, if $\mathbf{A} = \{l_1, l_3, l_3, l_5\}$, then $\nu(\mathbf{A}) = 3$ and $\phi_3(\mathbf{A}) = (0, 0, 1, 0, \dots, 0)$.

Figure 2: We estimate the projection-free 3D positions $\mathbf{P}_i$ and $\mathbf{P}_j$ considering that the persons are standing on the ground (left image). We then extract the angles θ_i and θ_j as well as the distance d_{ij} . The right image shows a top-view of the geometry between two persons and their associated variables.

4. Two-Person Mapping

As with the single-person case, we want to map the high-dimensional pairwise feature space into a low-dimensional space of posterior probabilities, that is,

$$\phi_2(\mathbf{y}_{ij}) = \left(P(\xi_{ij} = k | \mathbf{y}_{ij}) \right)_{k \in \mathcal{K}} \quad (4)$$

where $\mathcal{K}$ is a *pose-activity* dictionary and $\xi_{ij} \in \mathcal{K}$ defines the *pose-activity* word associated to feature $\mathbf{y}_{ij}$. However, the construction of such a dictionary is not straightforward as long as we need to learn *relevant* pose-activity words. Therefore, we learn this dictionary in an unsupervised manner using a spectral clustering method.

Suppose two persons are observed in a frame. Using the method of [13], we extract their 3D positions $\mathbf{P}_i$ and $\mathbf{P}_j$ as illustrated in Fig. 2. We define an interaction descriptor $\mathbf{p}_{ij}$ that encapsulates their relative distance and pose, as well as the associated activities: $\mathbf{p}_{ij} = (d_{ij}, \theta_i, \theta_j, a_i, a_j)$, where d_{ij} is the relative distance $d_{ij} = \|\mathbf{P}_i - \mathbf{P}_j\|_2$ and θ_i, θ_j are the angles spanned between the line $\mathbf{P}_i \mathbf{P}_j$ and the respective pose vectors. Fig. 2 (right) illustrates the geometry of the model. To learn the pose-activity words, we use a spectral clustering method [39] with a similarity function $\delta(\mathbf{p}_{ij}, \mathbf{p}_{kl}) = e^{-f(\mathbf{p}_{ij}, \mathbf{p}_{kl})^2}$, where f is defined by:

$$f(\mathbf{p}_{ij}, \mathbf{p}_{kl}) = \min \{ 2 - (\mathbb{1}(a_i = a_k) + \mathbb{1}(a_j = a_l)) + \lambda_1(\mu(\theta_i, \theta_k) + \mu(\theta_j, \theta_l)), 2 - (\mathbb{1}(a_i = a_l) + \mathbb{1}(a_j = a_k)) + \lambda_1(\mu(\theta_i, \theta_l) + \mu(\theta_j, \theta_k)) \} + \lambda_2 |d_{ij} - d_{kl}| \quad (5)$$

$$\text{with } \mu(\theta, \theta') = \min_{k \in \{1, 0, -1\}} |\theta - (\theta' + 2k\pi)|. \quad (6)$$

The parameters in (6), λ_1 and λ_2 are weights and the min operator is used to make the function f invariant to 2π shift-

ing and symmetric, that is, $f(p_{ij}, p_{kl})$ is equal to zero if $p_{ij} = p_{kl}$ or if $p_{ij} = p_{lk}$. The value of $f(p_{ij}, p_{kl})$ increases with the distance and the pose/action difference. As a consequence, the spectral clustering method tends to group together pairs of persons with the same activity, same distances and same relative poses. Considering a training set of interaction descriptors $\tilde{p}_{ij}$, the spectral clustering algorithm provides a cluster assignment label $\tilde{\xi}_{ij} \in \mathcal{K}$ for each descriptor. This label is used as a pairwise pose-activity word of the respective dictionary. As with single-person mapping, ϕ_2 is learned via a linear SVM. At test time, the assignment labels ξ_{ij} are not computed. Instead, ϕ_2 maps directly the high-dimensional representation onto the low-dimensional space where energy optimization is performed.

We illustrate our clustering approach on an example in order to clarify its contribution. For visualization simplicity we consider a dataset in which there is only one activity per image. As a consequence, $a_i = a_j$ holds for all interaction descriptors p_{ij} . After extracting each interaction descriptor, we apply the above described clustering method. Fig. 3 shows which cluster each interaction descriptor is assigned to. The interaction descriptors are displayed on one of the two plots according to the a_i value. In the case of *talking* activity, most of points are gathered in a cyan cluster where $(\theta_i, \theta_j) \approx (0, 0)$. This confirms the intuitive idea that the persons are likely to face each other when they are talking. On the contrary, the area $(\theta_i, \theta_j) \approx (0, 0)$ is almost empty in the case of the *queueing* activity since the persons are supposed to be in single-file when they are in a queue. Moreover, the only parameter that distinguishes the two clusters of the *queueing* activity is the distance d_{ij} .

5. Model Learning

We use structured support vector machine (SSVM) [25] to estimate the vectors $w_{a_i}^1, w_{a_i, a_j}^2$ and w^3 . To this end, Equation (3) is rewritten as

$$\begin{aligned} E(\mathbf{X}, \mathbf{Y}, \mathbf{A}) &= \sum_m^M w_{l_m}^{1\top} \left[\sum_{i=1}^I \phi_1(x_i) \mathbb{1}(l_m = a_i) \right] \\ &+ \sum_m^M w_{l_m}^{2\top} \left[\sum_{\substack{i,j \\ i \neq j}}^{I \times I} \phi_2(y_{ij}) \mathbb{1}(a_i = a_j) \mathbb{1}(l_m = a_i) \right] \\ &+ w^{3\top} \phi_3(\mathbf{A}) \\ &= w^\top \phi(\mathbf{X}, \mathbf{Y}, \mathbf{A}), \end{aligned} \quad (7)$$

Considering a set of N training images, From each of the $\tilde{I}^{(n)}$ bounding boxes, we extract the individual feature vectors $\tilde{\mathbf{X}}^{(n)}$ and the pairwise feature vectors $\tilde{\mathbf{Y}}^{(n)}$ to build

Figure 3: Clustering results obtained on dataset A' (see Sec.7 for more details about the datasets) for two activities: *talking*, *queueing*. Each color represents one cluster and each point is an interaction descriptor. θ_i and θ_j are measured in radians and d_{ij} in meters.

our training set $\{\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{A}}^{(n)}\}_{1 \leq n \leq N}$. We estimate w^* such that:

$$\begin{aligned} w^* &= \underset{w}{\operatorname{argmin}} \quad \frac{1}{2} \|w\|^2 \\ &+ C \sum_{n=1}^N \max_{A \in \mathcal{A}_{I(n)}} (w^\top \phi(\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \mathbf{A}) \\ &- w^\top \phi(\tilde{\mathbf{X}}^{(n)}, \tilde{\mathbf{Y}}^{(n)}, \tilde{\mathbf{A}}^{(n)}) + \Delta(\tilde{\mathbf{A}}^{(n)}, \mathbf{A})) \end{aligned} \quad (8)$$

Commonly, C is a hyper-parameter to adjust the SVM fitting, while $\Delta(\tilde{\mathbf{A}}^{(n)}, \mathbf{A})$ is the loss function that should penalize different labelings. Here, the following loss function

is used: $\Delta(\tilde{\mathbf{A}}^{(n)}, \mathbf{A}) = \sum_{i=1}^{\tilde{I}^{(n)}} \mathbb{1}(\tilde{a}_i^{(n)} \neq a_i)$. This function counts the number of incorrect labels. It is deliberately not

normalized to account the fact that the model can learn more from images with a large number of persons.

6. Feature Extraction

In this section we describe how we extract the individual (single-person) features $\mathbf{x}_i$ and the pairwise features $\mathbf{y}_{ij}$, that is the first-level features of our pipeline (see Fig. 1). The individual features $\mathbf{x}_i$ are built on two state-of-the-art feature extraction methods. The first one is based on point trajectories [37] and catches the local body motion whereas the second one is based on histogram of oriented gradients (HOG) [9] and catches the posture of a person. While local spatial gradients are described by the trajectories, HOG descriptor carries complementary information that is missing from trajectory features and is mainly useful for static activities. For each bounding box, we extract a trajectory-based feature vector $\mathbf{x}_i^{traj}$ and a HOG feature vector $\mathbf{x}_i^{hog}$. Two single-person mappings ϕ_1^{traj} and ϕ_1^{HOG} are learned. We then concatenate the two outputs (the low-dimensional vectors) to obtain $\phi_1(\mathbf{x}_i)$. Next, we explain how $\mathbf{x}_i^{traj}$ is built.

In order to use temporal information, we use the information from neighboring frames to build our box-wise descriptor. Thus, we propose to adapt the improved trajectory features of [37]. Improved trajectories consist of tracking over time densely sampled image points at different scales, while the local volume around the trajectory is encoded via several descriptors: HOG, histogram of optical flow, motion boundary histograms and trajectory shape, thus leading in a vector of dimension D . However, to build our single-person feature vector $\mathbf{x}_i^{traj}$, we proceed as follows. To describe a bounding-box of a person at frame t , we consider the interval $[t - \Delta t, t + \Delta t]$. First, sampled points at frame $t - \Delta t$ are tracked until frame t , and only trajectories that cross the bounding box of a person at frame t are kept. The remaining points are tracked until the frame $t + \Delta t$. Fig. 4 illustrates trajectories that are kept at frame t while their end points are further tracked. Recall that this allows us to use temporal information without building a complete temporal model. In [37], trajectories with low variance are rejected based on the assumption that they most likely belong to the background. Here, we keep all the trajectories since our region of interest (bounding box) does not contain lots of background information. Moreover, lack of motion seems to be informative for some activities like *waiting* or *queuing*. The dimension of the feature vectors is halved ($D/2$) via PCA. We then encode the individual features of each person into Fisher vectors of size κD using κ Gaussian components [33].

We also use trajectory-based features to obtain the pairwise (two persons) feature vectors $\mathbf{y}_{ij}$. Considering two persons (and their bounding boxes) detected in frame $t - \Delta t$,

Figure 4: Trajectory-based description of (pink) bounding-box at frame t : Points of frame $t - \Delta t$ are densely sampled and tracked up to frame t . At time t , (red) trajectories out of the bounding box are rejected. The remaining end points of trajectories (green) are further tracked until $t + \Delta t$. The local space-time area around the green trajectories is then encoded into the individual features $\mathbf{x}_i$.

we densely sample points and track them until frame t as with the individual features. Here, only trajectories that cross one of the bounding boxes at frame t are kept. As before, the remaining points are tracked until frame $t + \Delta t$. We encode the descriptors into Fisher vectors using κ Gaussian components to obtain the pairwise vector $\mathbf{y}_{ij}$.

7. Experiments

We first describe the datasets used to evaluate our model and give the implementation details. Choi *et al.* first published a collective activity dataset in [7]. This dataset is composed of 44 videos with 5 activities: *crossing*, *waiting*, *queueing*, *walking*, *talking*, and 8 poses: *right*, *front-right*, *front*, *front-left*, *left*, *back-left*, *back*, *back-right*. We refer here to this dataset as *dataset A*. This dataset is used to evaluate our model for the problem of recognizing multiple-group (distinct) activities. In order to compare with methods that can only deal with the single-group activity problem, two other datasets are used. Choi and Savarese published a new collective activity dataset [5], which we refer to as *dataset B*. This dataset consists in 33 videos with 6 activities: *gathering*, *talking*, *dismissal*, *walking together*, *chasing*, *queueing*. It is important to note that even if different activities can occur sequentially in a video of dataset *B*, all persons perform the same activity, which means that there is a single activity at each frame. Instead, dataset *A* may contain one or several distinct activities that are performed in parallel by different groups. At the same time, Choi and Savarese [5] published new annotations for dataset *A* in which they only consider one dominant activity per frame. In other words, if two groups of people perform two different activities, all the persons are considered as belonging to the activity of the largest group. We refer to this dataset as *dataset-A'*. As with dataset-*emphA*, eight pose labels are considered. Note that in both datasets *A'* and *B*, annotations are provided for one out of three frames, whereas one out of ten frames are annotated in dataset *A*.

As in [5], we split each dataset in training and test sets. One of the main difficulties with these datasets is that the classes are unbalanced. To figure out this issue, we use synthetic data augmentation as proposed in [3] and we produce new features that belong to the convex hull of the original features for minority classes, before training with SVM. While data augmentation does not directly apply to structured data, we augment the training examples of SSVM by perturbing the data of minor classes with Gaussian noise.

For the trajectory features, we use the values $D = 426$ and $\Delta t = 7$ in order to track points over 15 frames, as recommended in [37]. Then, Fisher vectors are built based on mixture of $\kappa = 128$ Gaussian components, while its dimensionality is reduced to 1500 using PCA, trained on a balanced set of features. As for pairwise features, we use weights $\lambda_1 = 0.33$ and $\lambda_2 = 0.25$. The weights have been chosen such that each term of the function f in (5) is equally important. The camera parameters used for the projection-free map estimation have been chosen to give results which look realistic but we cannot evaluate the predicted positions quantitatively because ground-truth values are not provided. However the satisfying clustering results, e.g. Fig. 3, confirm the quality of the 3D position estimations. We chose $K = 30$ for dataset A and $K = 12$ clusters for datasets A' and B .

7.1. Evaluation of the Proposed Model

In this section, we show results from a series of experiments in order to make clear the contribution of each component of the proposed methodological pipeline. We compare six variants of the proposed framework:

- *Single person*: Here, only the individual terms are used, namely the first double sum of (2). In the case of dataset A this is equivalent to a single SVM estimation per person. For dataset B , since we use the prior that all the persons are perform the same activity, our method yields better performance than single SVM estimations.
- *Static Pairwise*: Here we use simpler pairwise features on the following grounds. Instead of using the trajectories, we use HOG features and a linear SVM to estimate person poses, then we deduce an interaction descriptor $\mathbf{p}_{ij}$. We then estimate $(P(\xi_{ij} = k | \mathbf{p}_{ij}))_{k \in \mathcal{K}}$ with a second SVM. The posterior probabilities of each pose-activity word are then used as a two-person mapping ϕ_2 . This way, only one frame is used to extract the pairwise features.
- *Without regularization*: We use the model as described above but the regularization term in (2) is omitted.

Note that only the performance with dataset A is affected by this modification. As long as only one group activity occurs at a time in datasets A' and B , the regularizer has no impact.

- *Full model*: We make use of the full model, as proposed above.
- *4 meter threshold*: Since [19] observe some benefits from pruning, we test the full model along with the constraint that only person pairs whose distance is less than 4 meters are summed up in (3). We consider this threshold as a reasonable assumption for the used activities, i.e., people that are at least 4 meters away from each other are not taken into account.
- *Temporal smoothing*: A median filter is applied on the *full model* output to smooth predictions. Only the past predictions are used within a 12-frame window. As we choose $\Delta t = 7$ to extract the trajectories, a 26-frame window is used in total for each prediction. We test however the contribution of temporal smoothing only for datasets A' and B , since different activities occur simultaneously in dataset A and this would require to track all the persons.

	dataset A	dataset A'	dataset B
<i>Single person</i>	56.3%	59.6%	74.6%
<i>Static pairwise</i>	66.6%	73.8%	75.4%
<i>Without regularization</i>	73.3%	-	-
<i>Full model</i>	74.5%	79.8%	78.8%
<i>4 meter threshold</i>	74.8%	78.9%	78.4%
<i>Temporal smoothing</i>	-	81.5%	80.2%

Table 1: Average per-class activity recognition accuracy on three datasets for the single and multiple group activity problems

The performance of each one of these variants is evaluated on the three datasets. We summarize the results in Table 1. We can see that pairwise features drastically help the recognition in datasets A and A' . Fig. 5 illustrates this gain with a few examples. The gain is smaller for dataset B . This can be explained by the fact that dataset B constitutes an easier case compared to datasets A and A' : stable camera, simpler action scenarios, fewer persons, and less overlap. It validates the principle that taking pairwise interactions into consideration helps in complex and natural environments. Unlike [35, 19], the pruning strategy is not very helpful in our case, i.e., we do not consistently notice improvement when using a distance threshold. This may mean that other methods rely more on the used features compared to the proposed one. Our pairwise features succeed in capturing the information of the distance between persons, and

SSVM has learned that pairs with high distances have to be ignored. A temporal yet non-sophisticated smoother helps on dataset A' and B and even increases the performance of the proposed model.

7.2. Comparison with State of the Art Methods

Designing a benchmark for group activity recognition is difficult for the following reasons. Some authors use additional information/annotation for training or for testing, or they use different evaluation measures. For example, group labels are required to evaluate the method proposed in [35]. Table 2 and Table 3 show the performance of several methods for the single and multiple group activity problems. We use the per-class average to compute the performance considering the unbalanced number of available classes for each example. In the case of multiple activity recognition (Table 2) our method is the second-best performing method. Note however that additional group annotations are used by [35]. In the case of the recognition of single group activities (Table 3), our method also yields the second-best results. It should be noted that the best performing method, i.e. [1], uses its own person detections which biases the comparison with the other methods.

Choi <i>et al.</i> [7]	65.0%
Lan <i>et al.</i> [18]	68.2%
Sun <i>et al.</i> [35]	77.1%
Proposed	74.8%

Table 2: Classification accuracies based on *per-class averages* for multiple group-activity recognition on dataset A . Note that [35] uses additional annotations.

	dataset A'	dataset B	
Khamis <i>et al.</i> [16]	72.0%	-	
Lan <i>et al.</i> [20]	78.4%	-	#
Deng <i>et al.</i> [10]	80.6%	-	*
Ibrahim <i>et al.</i> [14]	80.9%	-	
Sun <i>et al.</i> [35]	81.2%	-	*
Hajimirsadeghi <i>et al.</i> [12]	81.9%	-	
Amer <i>et al.</i> [1]	92.2%	87.2%	*#
Choi & Savarese [6]	79.9%	79.2%	†
Nabi <i>et al.</i> [26]	-	72.4%	
Proposed	81.5%	80.2%	

* Only the average accuracies are provided.

Uses the authors' person detector.

† Uses extra annotations.

Table 3: Classification accuracies based on *per-class averages* for single group-activity recognition.

Fig. 6 plots the confusion matrices obtained with three methods and with our method. Note that it is not possi-

ble to show the confusion matrix obtained with [1] because only average results are provided in this paper. On dataset A , we compare our method with [35] which uses extra annotations. The very good recognition scores obtained with the proposed method for the *queueing* activity shows the added value brought in by the use of pairwise features. As expected, both methods have difficulties in discriminating between *crossing* and *walking*. It seems difficult to reduce the confusion between these two categories in the absence of context information. On dataset A' , the proposed method discriminates between *crossing* and *walking* slightly better than [6] which, again, uses extra annotations, namely interaction labels: *facing-each-other*, *standing-side-by-side*, and *standing-still*. In contrast, our method relies on two-person descriptors with no extra annotations.

On dataset B , the proposed method recognizes well activities for which motion plays an important role, such as *chasing* and *walking*. It performs less well whenever there are large motion variations during the activity, e.g. *gathering*. Nevertheless, our method yields a sensible result for this activity which is split into two parts: when the participants are spread out, the activity is labeled as *walking*, and once they get close to each other, it is labeled as *talking*. We believe that these are valid recognition results, although the activity is annotated as *gathering*. Nevertheless, activities with large dynamic variations might be dealt with by using a more sophisticated temporal model.

8. Conclusions

In this paper we proposed to model the group activity recognition problem as a structured labeling problem which is solved via an energy optimization framework. The proposed model can handle videos with either single or multiple group activities that occur simultaneously. Once high-dimensional feature vectors are calculated for each bounding-box and each for each pair of bounding boxes, we learn mapping functions in order to get meaningful representations in low-dimensional spaces. The geometry of the relationship between pairs of persons is included in the mapping function by learning a dictionary of relevant interactions. The geometry of the activity is thus taken into account without relying on pose estimations. We propose to use trajectory features to encode individual person description as well as the interactions between pairs of persons. The proposed framework outperforms several state-of-the-art methods for the recognition of multiple group-activities and compares favorably with other methods for the recognition of a single group activity. Future work includes the further use of group features as well as adding a temporal model to the energy formulation.

Figure 5: Example frames from sequences of dataset *A* with single person model (first row) and the proposed model (second row) . The color of the box represents the predicted activity (red for *Crossing*, yellow for *Waiting*, green for *Queueing*, blue for *Walking*, purple for *Talking*). On top of each bounding box, the letters show the ground truth activity (*better viewed on screen*)

Figure 6: Confusion matrices for dataset *A* (left column), dataset *A'* (middle column) and dataset *B* (right column).

References

- [1] M. R. Amer, P. Lei, and S. Todorovic. HIRF: Hierarchical random field for collective activity recognition in videos. In *ECCV*, 2014.
- [2] M. R. Amer and S. Todorovic. A chains model for localizing participants of group activities in videos. In *ICCV*, 2011.
- [3] N. Chawla, K. Bowyer, L. Hall, and W. Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *JAIR*, 2002.
- [4] W. Choi, Y.-W. Chao, C. Pantofaru, and S. Savarese. Discovering groups of people in images. In *ECCV*, 2014.
- [5] W. Choi and S. Savarese. A unified framework for multi-target tracking and collective activity recognition. In *ECCV*, 2012.
- [6] W. Choi and S. Savarese. Understanding collective activities of people from videos. *IEEE TPAMI*, 2013.
- [7] W. Choi, K. Shahid, and S. Savarese. What are they doing?: Collective activity classification using spatio-temporal relationship among people. In *ICCV Workshops*, 2009.
- [8] F. Cupillard, F. Brmond, and M. Thonnat. Group behavior recognition with multiple cameras. In *WACV*, 2002.
- [9] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005.
- [10] Z. Deng, M. Zhai, L. Chen, Y. Liu, S. Muralidharan, M. J. Roshtkhari, and G. Mori. Deep structured models for group activity recognition. In *BMVC*, 2015.
- [11] A. Gupta, A. Kembhavi, and L. S. Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE TPAMI*, 2009.
- [12] H. Hajimirsadeghi, W. Yan, A. Vahdat, and G. Mori. Visual recognition by counting instances: A multi-instance cardinality potential kernel. *arXiv preprint arXiv:1502.02063*, 2015.
- [13] D. Hoiem, A. A. Efros, and M. Hebert. Putting objects in perspective. *IJCV*, 2008.
- [14] M. Ibrahim, S. Muralidharan, Z. Deng, A. Vahdat, and G. Mori. A hierarchical deep temporal model for group activity recognition. In *CVPR*, 2016.
- [15] M. Jain, J. C. van Gemert, and C. G. M. Snoek. What do 15,000 object categories tell us about classifying and localizing actions? In *CVPR*, 2015.
- [16] S. Khamis, V. I. Morariu, and L. S. Davis. Combining per-frame and per-track cues for multi-person action recognition. In *ECCV*, 2012.
- [17] H. Kjellstrm, J. Romero, D. Martnez, and D. Kragi. Simultaneous visual recognition of manipulation actions and manipulated objects. In *ECCV*, 2008.
- [18] T. Lan, Y. Wang, G. Mori, and S. N. Robinovitch. Retrieving actions in group contexts. In *Trends and Topics in Computer Vision*, 2010.
- [19] T. Lan, Y. Wang, W. Yang, and G. Mori. Beyond actions: Discriminative models for contextual group activities. In *NIPS*, 2010.
- [20] T. Lan, Y. Wang, W. Yang, S. N. Robinovitch, and G. Mori. Discriminative latent models for recognizing contextual group activities. *IEEE TPAMI*, 2012.
- [21] Z. Lan, M. Lin, X. Li, A. G. Hauptmann, and B. Raj. Beyond Gaussian pyramid: Multi-skip feature stacking for action recognition. In *CVPR*, 2015.
- [22] R. Li, R. Chellappa, and S. K. Zhou. Learning multi-modal densities on discriminative temporal interaction manifold for group activity recognition. In *CVPR*, 2009.
- [23] M. Marszalek, I. Laptev, and C. Schmid. Actions in context. In *CVPR*, 2009.
- [24] K. Murphy, A. Torralba, and W. Freeman. Using the forest to see the trees: a graphical model relating features, objects and scenes. In *NIPS*, 2003.
- [25] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [26] M. Nabi, A. Del Bue, and V. Murino. Temporal poselets for collective activity detection and recognition. In *ICCVW Workshops*, 2013.
- [27] S. Odashima, M. Shimosaka, T. Kaneko, R. Fukui, and T. Sato. Collective activity localization with contextual spatial pyramid. In *ECCV*, 2012.
- [28] R. Poppe. A survey on vision-based human action recognition. *IVC*, 2010.
- [29] A. Rabinovich, A. Vedaldi, C. Galleguillos, E. Wiewiora, and S. Belongie. Objects in context. In *ICCV*, 2007.
- [30] M. Ryoo and J. Aggarwal. Stochastic representation and recognition of high-level group activities. *IJCV*, 2011.
- [31] M. S. Ryoo and J. K. Aggarwal. Recognition of composite human activities through context-free grammar based representation. In *CVPR*, 2006.
- [32] M. S. Ryoo and J. K. Aggarwal. Spatio-temporal relationship match: Video structure comparison for recognition of complex human activities. In *ICCV*, 2009.
- [33] J. Snchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. *IJCV*, 2013.
- [34] C. Schuldt, I. Laptev, and B. Caputo. Recognizing human actions: a local SVM approach. In *ICPR*, 2004.
- [35] L. Sun, H. Ai, and S. Lao. Activity group localization by modeling the relations among participants. In *ECCV*, 2014.
- [36] P. Turaga, R. Chellappa, V. S. Subrahmanian, and O. Udrea. Machine recognition of human activities: A survey. *IEEE TCSVT*, 2008.
- [37] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [38] B. Yao and L. Fei-Fei. Recognizing human-object interactions in still images by modeling the mutual context of objects and human poses. *IEEE TPAMI*, 2012.
- [39] S. X. Yu and J. Shi. Multiclass spectral clustering. In *ICCV*, 2003.