

PoseTrack: Joint Multi-Person Pose Estimation and Tracking

Umar Iqbal¹, Anton Milan², and Juergen Gall¹

¹Computer Vision Group, University of Bonn, Germany

²Australian Centre for Visual Technologies, University of Adelaide, Australia

Abstract

In this work, we introduce the challenging problem of joint multi-person pose estimation and tracking of an unknown number of persons in unconstrained videos. Existing methods for multi-person pose estimation in images cannot be applied directly to this problem, since it also requires to solve the problem of person association over time in addition to the pose estimation for each person. We therefore propose a novel method that jointly models multi-person pose estimation and tracking in a single formulation. To this end, we represent body joint detections in a video by a spatio-temporal graph and solve an integer linear program to partition the graph into sub-graphs that correspond to plausible body pose trajectories for each person. The proposed approach implicitly handles occlusion and truncation of persons. Since the problem has not been addressed quantitatively in the literature, we introduce a challenging “Multi-Person PoseTrack” dataset, and also propose a completely unconstrained evaluation protocol that does not make any assumptions about the scale, size, location or the number of persons. Finally, we evaluate the proposed approach and several baseline methods on our new dataset.

1. Introduction

Human pose estimation has long been motivated for its applications in understanding human interactions, activity recognition, video surveillance and sports video analytics. The field of human pose estimation in images has progressed remarkably over the past few years. The methods have advanced from pose estimation of single pre-localized persons [30, 6, 40, 14, 16, 27, 5, 32] to the more challenging and realistic case of multiple, potentially overlapping and truncated persons [12, 8, 30, 16, 17]. Many applications, such as mentioned before, however, aim to analyze human body motion over time. While there exists a notable number of works that track the pose of a single person in a video [28, 9, 44, 33, 46, 20, 29, 7, 13, 18], multi-person human pose estimation in unconstrained videos has not been addressed in the literature.




Figure 1: Example frames and annotations from the proposed Multi-Person PoseTrack dataset.

In this work, we address the problem of tracking the poses of multiple persons in an unconstrained setting. This means that we have to deal with large pose and scale variations, fast motions, and a varying number of persons and visible body parts due to occlusion or truncation. In contrast to previous works, we aim to solve the association of each person across the video and the pose estimation together. To this end, we build upon the recent methods for multi-person pose estimation in images [30, 16, 17] that build a spatial graph based on joint proposals to estimate the pose of multiple persons in an image. In particular, we cast the problem as an optimization of a densely connected spatio-temporal graph connecting body joint candidates spatially as well as temporally. The optimization problem is formulated as a constrained Integer Linear Program (ILP) whose feasible solution partitions the graph into valid body pose trajectories for any unknown number of persons. In this way, we can handle occlusion, truncation, and temporal association within a single formulation.

Since there exists no dataset that provides annotations to quantitatively evaluate joint multi-person pose estimation and tracking, we also propose a new challenging *Multi-Person PoseTrack* dataset as a second contribution of the paper. The dataset provides detailed and dense annotations for multiple persons in each video, as shown in Fig. 1, and introduces new challenges to the field of pose estimation in videos. In order to evaluate the pose estimation and tracking accuracy, we introduce a new protocol that also deals with occluded body joints. We quantify the proposed method in detail on the proposed dataset, and also report results for several baseline methods. The source code, pre-trained models and the dataset are publicly available.¹

¹http://pages.iai.uni-bonn.de/iqbal_umar/PoseTrack/

2. Related Work

Single person pose estimation in images has seen a remarkable progress over the past few years [39, 30, 6, 40, 14, 16, 27, 5, 32]. However, all these approaches assume that only a single person is visible in the image, and cannot handle realistic cases where several people appear in the scene, and interact with each other. In contrast to single person pose estimation, multi-person pose estimation introduces significantly more challenges, since the number of persons in an image is not known a priori. Moreover, it is natural that persons occlude each other during interactions, and may also become partially truncated to various degrees. Multi-person pose estimation has therefore gained much attention recently [11, 37, 31, 43, 23, 12, 8, 3, 30, 16, 17]. Earlier methods in this direction follow a two-staged approach [31, 12, 8] by first detecting the persons in an image followed by a human pose estimation technique for each person individually. Such approaches are, however, applicable only if people appear well separated and do not occlude each other. Moreover, most single person pose estimation methods always output a fixed number of body joints and do not account for occlusion and truncation, which often is the case in multi-person scenarios. Other approaches address the problem using tree structured graphical models [43, 37, 11, 23]. However, such models struggle to cope with large pose variations, and are shown to be significantly outperformed by more recent methods based on Convolutional Neural Networks [30, 16]. For example, [30] jointly estimate the pose of all persons visible in an image, while also handling occlusion and truncation. The approach has been further improved by stronger part detectors and efficient approximations [16]. The approach in [17] also proposes a simplification of [30] by tackling the problem locally for each person. However, it still relies on a separate person detector.

Single person pose estimation in videos has also been studied extensively in the literature [28, 9, 46, 33, 46, 20, 44, 29, 13, 18]. These approaches mainly aim to improve pose estimation by utilizing temporal smoothing constraints [28, 9, 44, 33, 13] and/or optical flow information [46, 20, 29], but they are not directly applicable to videos with multiple potentially occluding persons.

In this work we focus on the challenging problem of joint multi-person pose estimation and data association across frames. While the problem has not been studied quantitatively in the literature², there exist early works towards the problem [19, 2]. These approaches, however, do not reason jointly about pose estimation and tracking, but rather focus on multi-person tracking alone. The methods follow a multi-staged strategy, *i.e.* they first estimate body part loca-

²Contemporaneously with this work, the problem has also been studied in [15]

tions for each person separately and subsequently leverage body part tracklets to facilitate person tracking. We on the other hand propose to simultaneously estimate the pose of multiple persons and track them over time. To this end, we build upon the recent progress on multi-person pose estimation in images [30, 16, 17] and propose a joint objective for both problems.

Previous datasets used to benchmark pose estimation algorithms in-the-wild are summarized in Tab. 1. While there exists a number of datasets to evaluate single person pose estimation methods in videos, such as *e.g.*, J-HMDB [21] and Penn-Action [45], none of the video datasets provides annotations to benchmark multi-person pose estimation and tracking at the same time. To allow for a quantitative evaluation of this problem, we therefore also introduce a new “Multi-Person PoseTrack” dataset which provides pose annotations for multiple persons in each video to measure pose estimation accuracy, and also provides a unique ID for each of the annotated persons to benchmark multi-person pose tracking. The proposed dataset introduces new challenges to the field of human pose estimation and tracking since it contains a large amount of appearance and pose variations, body part occlusion and truncation, large scale variations, fast camera and person movements, motion blur, and a sufficiently large number of persons per video.

3. Multi-Person Pose Tracking

Our method jointly solves the problem of multi-person pose estimation and tracking for all persons appearing in a video together. We first generate a set of joint detection candidates in each video as illustrated in Fig. 2. From the detections, we build a graph consisting of spatial edges connecting the detections within a frame and temporal edges connecting detections of the same joint type over frames. We solve the problem using integer linear programming (ILP) whose feasible solution provides the pose estimate for each person in all video frames, and also performs person association across frames. We first introduce the proposed method and discuss the proposed dataset for evaluation in Sec. 4.

3.1. Spatio-Temporal Graph

Given a video sequence \mathcal{F} containing an arbitrary number of persons, we generate a set of body joint detection candidates $D = \{D_f\}_{f \in \mathcal{F}}$ where D_f is the set for frame f . Every detection $d \in D$ at location $\mathbf{x}_d^f \in \mathbb{R}^2$ in frame f belongs to a joint type $j \in \mathcal{J} = \{1, \dots, J\}$. Additional details regarding the used detector will be provided in Sec. 3.4.

For multi-person pose tracking, we aim to identify the joint hypotheses that belong to an individual person in the entire video. This can be formulated by a graph structure $G = (D, E)$ where D is the set of nodes. The set of edges E consists of two types of edges, namely spatial edges E_s and temporal edges E_t . The spatial edges correspond to the




Figure 2: **Top:** Body joint detection hypotheses shown for three frames. **Middle:** Spatio-temporal graph with spatial edges (blue) and temporal edges for head (red) and neck (yellow). We only show a subset of the edges. **Bottom:** Estimated poses for all persons in the video. Each color corresponds to a unique person identity.

union of edges of a fully connected graph for each frame, *i.e.*

$$E_s = \bigcup_{f \in \mathcal{F}} E_s^f \text{ and } E_s^f = \{(d, d') : d \neq d' \wedge d, d' \in D_f\}. \quad (1)$$

Note that these edges connect joint candidates independently of the associated joint type j . The temporal edges connect only joint hypotheses of the same joint type over two different frames, *i.e.*

$$E_t = \{(d, d') : j=j' \wedge d \in D_f \wedge d' \in D_{f'} \wedge 1 \leq |f - f'| \leq \tau \wedge f, f' \in \mathcal{F}\}. \quad (2)$$

The temporal connections are not only modeled for neighboring frames, *i.e.* $|f - f'| = 1$, but we also take temporal relations up to τ frames into account to handle short-term occlusion and missing detections. The graph structure is illustrated in Fig. 2.

3.2. Graph Partitioning

By removing edges and nodes from the graph $G = (D, E)$, we obtain several partitions of the spatio-temporal graph and each partition corresponds to a tracked pose of an individual person. In order to solve the graph partitioning problem, we introduce the three binary vectors $v \in \{0, 1\}^{|D|}$, $s \in \{0, 1\}^{|E_s|}$, and $t \in \{0, 1\}^{|E_t|}$. Each binary variable implies if a node or edge is removed, *i.e.* $v_d=0$ implies that the joint detection d is removed. Similarly, $s_{(d_f, d'_{f'})}=0$ with $(d_f, d'_{f'}) \in E_s$ implies that the spatial edge between the joint hypothesis d and d' in frame f

is removed while $t_{(d_f, d'_{f'})}=0$ with $(d_f, d'_{f'}) \in E_t$ implies that the temporal edge between the joint hypothesis d in frame f and d' in frame f' is removed.

A partitioning is obtained by minimizing the cost function

$$\operatorname{argmin}_{v, s, t} (\langle v, \phi \rangle + \langle s, \psi_s \rangle + \langle t, \psi_t \rangle) \quad (3)$$

$$\langle v, \phi \rangle = \sum_{d \in D} v_d \phi(d) \quad (4)$$

$$\langle s, \psi_s \rangle = \sum_{(d_f, d'_{f'}) \in E_s} s_{(d_f, d'_{f'})} \psi_s(d_f, d'_{f'}) \quad (5)$$

$$\langle t, \psi_t \rangle = \sum_{(d_f, d'_{f'}) \in E_t} t_{(d_f, d'_{f'})} \psi_t(d_f, d'_{f'}). \quad (6)$$

This means that we search for a graph partitioning such that the cost of the remaining nodes and edges is minimal. The cost for a node d is defined by the unary term:

$$\phi(d) = \log \frac{1 - p_d}{p_d} \quad (7)$$

where $p_d \in (0, 1)$ corresponds to the probability of the joint hypothesis d . Note that $\phi(d)$ is negative when $p_d > 0.5$ and detections with a high confidence are preferred since they reduce the cost function (3). The cost for a spatial or tem-

poral edge is defined similarly by

$$\psi_s(d_f, d'_f) = \log \frac{1 - p^s_{(d_f, d'_f)}}{p^s_{(d_f, d'_f)}} \quad (8)$$

$$\psi_t(d_f, d'_{f'}) = \log \frac{1 - p^t_{(d_f, d'_{f'})}}{p^t_{(d_f, d'_{f'})}}. \quad (9)$$

While p^s denotes the probability that two joint detections d and d' in a frame f belong to the same person, p^t denotes the probability that two detections of a joint in frame f and f' are the same. In Sec. 3.4 we will discuss how the probabilities p_d , $p^s_{(d_f, d'_f)}$, and $p^t_{(d_f, d'_{f'})}$ are learned.

In order to ensure that the feasible solutions of the objective (3) result in well defined body poses and valid pose tracks, we have to add additional constraints. The first set of constraints ensures that two joint hypotheses are associated to the same person ($s_{(d_f, d'_f)}=1$) only if both detections are considered as valid, *i.e.*, $v_{d_f}=1$ and $v_{d'_f}=1$:

$$s_{(d_f, d'_f)} \leq v_{d_f} \wedge s_{(d_f, d'_f)} \leq v_{d'_f} \quad \forall (d_f, d'_f) \in E_s. \quad (10)$$

The same holds for the temporal edges:

$$t_{(d_f, d'_{f'})} \leq v_{d_f} \wedge t_{(d_f, d'_{f'})} \leq v_{d'_{f'}} \quad \forall (d_f, d'_{f'}) \in E_t. \quad (11)$$

The second set of constraints are transitivity constraints in the spatial domain. Such transitivity constraints have been proposed for multi-person pose estimation in images [30, 16, 17]. They enforce for any triplet of joint detection candidates (d_f, d'_f, d''_f) that if d_f and d'_f are associated to one person and d'_f and d''_f are also associated to one person, *i.e.* $s_{(d_f, d'_f)}=1$ and $s_{(d'_f, d''_f)}=1$, then the edge (d_f, d''_f) should also be added:

$$s_{(d_f, d'_f)} + s_{(d'_f, d''_f)} - 1 \leq s_{(d_f, d''_f)} \quad (12)$$

$$\forall (d_f, d'_f), (d'_f, d''_f) \in E_s.$$

An example of a transitivity constraint is illustrated in Fig. 3a. The transitivity constraints can be used to enforce that a human can have only one joint type j , *e.g.* only one head. Let d_f and d''_f have the same joint type j while d'_f belongs to another joint type j' . Without transitivity constraints connecting d_f and d''_f with d'_f might result in a low cost. The transitivity constraints, however, enforce that the binary cost $\psi_s(d_f, d''_f)$ is added. To prevent poses with multiple joints, we thus only have to ensure that the binary cost $\psi_s(d, d'')$ is very high if $j=j''$. We discuss this more in detail in Sec. 3.4.

In contrast to previous work, we also have to ensure spatio-temporal consistency. Similar to the spatial transitivity constraints (12), we can define temporal transitivity constraints:

$$t_{(d_f, d'_{f'})} + t_{(d'_{f'}, d''_{f''})} - 1 \leq t_{(d_f, d''_{f''})} \quad (13)$$

$$\forall (d_f, d'_{f'}), (d'_{f'}, d''_{f''}) \in E_t.$$




Figure 3: (a) The spatial transitivity constraints (12) ensure that if the two joint hypotheses d_f and d''_f are spatially connected to d'_f (red edges) then the cost of the spatial edge between d_f and d''_f (green edge) also has to be added. (b) The temporal transitivity constraints (13) ensure transitivity for temporal edges (dashed). (c) The spatio-temporal transitivity constraints (14) model transitivity for two temporal edges and one spatial edge. (d) The spatio-temporal consistency constraints (15) ensure that if two pairs of joint hypotheses $(d_f, d'_{f'})$ and $(d''_{f'}, d'''_{f''})$ are temporally connected (dashed red edges) and d_f and d''_f are spatially connected (solid red edge) then the cost of the spatial edge between $d'_{f'}$ and $d'''_{f''}$ (solid green edge) also has to be added.

The last set of constraints are spatio-temporal constraints that ensure that the pose is consistent over time. We define two types of spatio-temporal constraints. The first type consists of a triplet of joint detection candidates $(d_f, d'_{f'}, d''_{f''})$ from two different frames f and f' . The constraints are defined as,

$$t_{(d_f, d'_{f'})} + t_{(d_f, d''_{f''})} - 1 \leq s_{(d'_{f'}, d''_{f''})}$$

$$t_{(d_f, d'_{f'})} + s_{(d'_{f'}, d''_{f''})} - 1 \leq t_{(d_f, d''_{f''})} \quad (14)$$

$$\forall (d_f, d'_{f'}), (d_f, d''_{f''}) \in E_t,$$

and enforce transitivity for two temporal edges and one spatial edge. The second type of spatio-temporal constraints are based on quadruples of joint detection candidates $(d_f, d'_{f'}, d''_{f''}, d'''_{f'''})$ from two different frames f and f' . The spatio-temporal constraints ensure that if $(d_f, d'_{f'})$ and $(d''_{f''}, d'''_{f'''})$ are temporally connected and $(d_f, d''_{f''})$ are spatially connected then the spatial edge $(d'_{f'}, d'''_{f'''})$ has to be added:

$$t_{(d_f, d'_{f'})} + t_{(d''_{f''}, d'''_{f'''})} + s_{(d_f, d''_{f''})} - 2 \leq s_{(d'_{f'}, d'''_{f'''})}$$

$$t_{(d_f, d'_{f'})} + t_{(d''_{f''}, d'''_{f'''})} + s_{(d'_{f'}, d''_{f''})} - 2 \leq s_{(d_f, d''_{f''})} \quad (15)$$

$$\forall (d_f, d'_{f'}), (d''_{f''}, d'''_{f'''}) \in E_t.$$

An example of both types of spatio-temporal constraint can be seen in Fig. 3c and Fig. 3d, respectively.

3.3. Optimization

We optimize the objective (3) with the branch-and-cut algorithm of the ILP solver Gurobi. To reduce the runtime for long sequences, we process the video batch-wise where each batch consists of $k = 31$ frames. For the first k frames, we build the spatio-temporal graph as discussed and optimize the objective (3). We then continue to build a graph

for the next k frames and add the previously selected nodes and edges to the graph, but fix them such that they cannot be removed anymore. Since the graph partitioning produces also small partitions, which usually correspond to clusters of false positive joint detections, we remove any partition that is shorter than 7 frames or has less than 6 nodes per frame on average.

3.4. Potentials

In order to compute the unaries ϕ (7) and binaries ψ (8),(9), we have to learn the probabilities p_d , $p_{(d_f, d'_f)}$, and $p_{(d_f, d'_f)'}^t$.

The probability p_d is given by the confidence of the joint detector. As joint detector, we use the publicly available pre-trained CNN [16] trained on the MPII Multi-Person Pose dataset [30]. In contrast to [16], we do not assume that any scale information is given. We therefore apply the detector to an image pyramid with 4 scales $\gamma \in \{0.6, 0.9, 1.2, 1.5\}$. For each detection d located at \mathbf{x}_d^f , we compute a quadratic bounding box $B_d = \{\mathbf{x}_d^f, h_d\}$. We use $h_d = \frac{70}{\gamma}$ for the width and height. To reduce the number of detections, we remove all bounding boxes that have an intersection-over-union (IoU) ratio over 0.7 with another bounding box that has a higher detection confidence.

The spatial probability $p_{(d_f, d'_f)}^s$ depends on the joint types j and j' of the detections. If $j=j'$, we define $p_{(d_f, d'_f)}^s = \text{IoU}(B_d, B_{d'})$. This means that a joint type j cannot be added multiple times to a person except if the detections are very close. If a partition includes detections of the same type in a single frame, the detections are merged by computing the weighted mean of the detections, where the weights are proportional to p_d . If $j \neq j'$, we use the pre-trained binaries [16] after a scale normalization.

The temporal probability $p_{(d_f, d'_f)'}^t$ should be high if two detections of the same joint type at different frames belong to the same person. To that end, we build on the idea recently used in multi-person tracking [38] and compute dense correspondences between two frames using DeepMatching [41]. Let K_{d_f} and $K_{d'_f}$ be the sets of matched key-points inside the bounding boxes B_{d_f} and $B_{d'_f}$ and $\underline{K}_{dd'} = |K_{d_f} \cup K_{d'_f}|$ and $\overline{K}_{dd'} = |K_{d_f} \cap K_{d'_f}|$ the union and intersection of these two sets. We then form a feature vector by $\{\overline{K}/\underline{K}, \min(p_d, p_{d'}), \Delta \mathbf{x}_{dd'}, \|\Delta \mathbf{x}_{dd'}\|\}$ where $\Delta \mathbf{x}_{dd'} = \mathbf{x}_d^f - \mathbf{x}_{d'}^{f'}$. We also append the feature vector with non-linear terms as done in [38]. The mapping from the feature vector to the probability $p_{(d_f, d'_f)'}^t$ is obtained by logistic regression.

4. The Multi-Person PoseTrack Dataset

In this section we introduce our new dataset for multi-person pose estimation in videos. The MPII Multi-Person Pose [1] is currently one of the most popular benchmarks

Dataset	Video-labeled poses	multi-person	Large scale variation	variable skeleton size	# of persons
Leeds Sports [22]					2000
MPII Pose [1]			✓	✓	40,522
We Are Family [11]		✓			3131
MPII Multi-Person Pose [30]		✓	✓	✓	14,161
MS-COCO Keypoints [25]		✓	✓	✓	105,698
J-HMDB [21]	✓		✓	✓	32,173
Penn-Action [45]	✓		✓		159,633
VideoPose [35]	✓				1286
Poses-in-the-wild [9]	✓				831
YouTube Pose [7]	✓				5000
FYDP [36]	✓				1680
UYDP [36]	✓				2000
Multi-Person PoseTrack	✓	✓	✓	✓	16,219

Table 1: A comparison of PoseTrack dataset with the existing related datasets for human pose estimation in images and videos.

for multi-person pose estimation in images, and covers a wide range of activities. For each annotated image, the dataset also provides unlabeled video clips ranging 20 frames both forward and backward in time relative to that image. For our video dataset, we manually select a subset of all available videos that contain multiple persons and cover a wide variety of person-person or person-object interactions. Moreover, the selected videos are chosen to contain a large amount of body pose appearance and scale variation, as well as body part occlusion and truncation. The videos also contain severe body motion, *i.e.*, people occlude each other, re-appear after complete occlusion, vary in scale across the video, and also significantly change their body pose. The number of visible persons and body parts may also vary during the video. The duration of all provided video clips is exactly 41 frames. To include longer and variable-length sequences, we downloaded the original raw video clips using the provided URLs and obtained an additional set of videos. To prevent an overlap with the existing data, we only considered sequences that are at least 150 frames apart from the training samples, and followed the same rationale as above to ensure diversity.

In total, we compiled a set of 60 videos with the number of frames per video ranging between 41 and 151. The number of persons ranges between 2 and 16 with an average of more than 5 persons per video sequence, totaling over 16,000 annotated poses. The person heights are between 100 and 1200 pixels. We split the dataset into a training and testing set with an equal number of videos.

4.1. Annotation

As in [1], we annotate 14 body joints and a rectangle enclosing the person’s head. The latter is required to estimate the absolute scale which is used for evaluation. We assign a unique identity to every person appearing in the video. This person ID remains the same throughout the video un-

til the person moves out of the field-of-view. Since we do not target person re-identification in this work, we assign a new ID if a person re-appears in the video. We also provide occlusion flags for all body joints. A joint is marked occluded if it was in the field-of-view but became invisible due to an occlusion. Truncated joints, *i.e.* those outside the image border limits, are not annotated, therefore, the number of joints per person varies across the dataset. Very small persons were zoomed in to a reasonable size to accurately perform the annotation. To ensure a high quality of the annotation, all annotations were performed by trained in-house workers, following a clearly defined protocol. An example annotation can be seen in Fig. 1.

4.2. Experimental setup and evaluation metrics

Since the problem of simultaneous multi-person pose estimation and person tracking has not been quantitatively evaluated in the literature, we define a new evaluation protocol for this problem. To this end, we follow the best practices followed in both multi-person pose estimation [30] and multi-target tracking [26]. In order to evaluate whether a part is predicted correctly, we use the widely adopted PCKh (head-normalized probability of correct keypoint) metric [1], which considers a body joint to be correctly localized if the predicted location of the joint is within a certain threshold from the true location. Due to the large scale variation of people across videos and even within a frame, this threshold needs to be selected adaptively, based on the person’s size. To that end, [1] propose to use 30% of the head box diagonal. We have found this threshold to be too relaxed because recent pose estimation approaches are capable of predicting the joint locations rather accurately. Therefore, we use a more strict evaluation with a 20% threshold.

Given the joint localization threshold for each person, we compute two sets of evaluation metrics, one adopted from the multi-target tracking literature [42, 10, 26] to evaluate multi-person pose tracking, and one which is commonly used for evaluating multi-person pose estimation [30].

Tracking. To evaluate multi-person pose tracking, we consider each joint trajectory as one individual target,³ and compute multiple measures. First, the CLEAR MOT metrics [4] provide the tracking accuracy (MOTA) and tracking precision (MOTP). The former is derived from three types of error ratios: false positives, missed targets, and identity switches (IDs). These are linearly combined to produce a normalized accuracy where 100% corresponds to zero errors. MOTP measures how precise each object, or in our case each body joint, has been localized w.r.t. the ground-truth. Second, we report trajectory-based measures proposed in [24], that count the number of mostly tracked (MT) and mostly lost (ML) tracks. A track is considered mostly tracked if it has been recovered in at least 80% of its length,

³Note that only joints of the same type are matched.

and mostly lost if more than 80% are not tracked. For completeness, we also compute the number of times a ground-truth trajectory is fragmented (FM).

Pose. For measuring frame-wise multi-person pose accuracy, we use *Mean Average Precision* (mAP) as is done in [30]. The protocol to evaluate multi-person pose estimation in [30] assumes that the rough scale and location of a group of persons is known during testing [30], which is not the case in realistic scenarios, and in particular in videos. We therefore propose to make no assumption during testing and evaluate the predictions without rescaling or shifting them according to the ground-truth.

Occlusion handling. Both of the aforementioned protocols to measure pose estimation and tracking accuracy do not consider occlusion during evaluation, and penalize if an occluded target that is annotated in the ground-truth is not correctly estimated [26, 30]. This, however, discourages methods that either detect occlusion and do not predict the occluded joints or approaches that predict the joint position even for occluded joints. We want to provide a fair comparison for both types of occlusion handling. We therefore extend both measures to incorporate occlusion information explicitly. To this end, we first assign each person to one of the ground-truth poses based on the PCKh measure as done in [30]. For each matched person, we consider an occluded joint correctly estimated either if *a*) it is predicted at the correct location despite being occluded, or *b*) it is not predicted at all. Otherwise, the prediction is considered as a false positive.

5. Experiments

In this section we evaluate the proposed method for joint multi-person pose estimation and tracking on the newly introduced Multi-Person PoseTrack dataset.

5.1. Multi-Person Pose Tracking

The results for multi-person pose tracking (MOT CLEAR metrics) are reported in Tab. 2. To find the best setting, we first perform a series of experiments, investigating the influence of temporal connection density, temporal connection length, and inclusion of different constraint types.

We first examine the impact of different joint combinations for temporal connections. Connecting only the Head Tops (HT) between frames results in a Multi-Object Tracking Accuracy (MOTA) of 27.2 with a recall and precision of 57.6% and 66.0%, respectively. Adding Neck and Shoulder (HT:N:S) detections for temporal connections improves the MOTA score to 28.2, while also improving the recall from 57.6% to 62.7%. Adding more temporal connections also increases other metrics such as MT, ML, and also results in a lower number of ID switches (IDs) and fragments (FM). However, increasing the number of joints for temporal edges even further (HT:N:S:H) results in a slight de-

Method	RcII ↑	Prcn ↑	MT ↑	ML ↓	IDs ↓	FM ↓	MOTA ↑	MOTP ↑
<i>Impact of temporal connection density</i>								
HT	57.6	66.0	632	623	674	5080	27.2	56.1
HT:N:S	62.7	64.9	760	510	470	5557	28.2	55.8
HT:N:S:H	63.1	64.5	774	494	478	5564	27.8	55.7
HT:W:A	62.8	64.9	758	526	516	5458	28.2	55.8
<i>Impact of the length of temporal connection (τ)</i>								
HT:N:S ($\tau = 1$)	62.7	64.9	760	510	470	5557	28.2	55.8
HT:N:S ($\tau = 3$)	63.0	64.8	775	502	431	5629	28.2	55.7
HT:N:S ($\tau = 5$)	62.8	64.7	763	508	381	5676	28.0	55.7
<i>Impact of the constraints</i>								
All	63.0	64.8	775	502	431	5629	28.2	55.7
All \ spat. transitivity	22.2	76.0	115	1521	39	3947	15.1	58.0
All \ temp. transitivity	60.3	65.1	712	544	268	5610	27.7	55.8
All \ spatio-temporal	55.1	64.1	592	628	262	5444	23.9	55.7
<i>Comparison with the Baselines</i>								
Ours	63.0	64.8	775	502	431	5629	28.2	55.7
BBox-Tracking [38, 34]								
+ LJPA [17]	58.8	64.8	716	646	319	5026	26.6	53.5
+ CPM [40]	60.1	57.7	754	611	347	4969	15.6	53.4

Table 2: Quantitative evaluation of multi-person pose-tracking using common multi-object tracking metrics. Up and down arrows indicate whether higher or lower values for each metric are better. The first three blocks of the table present an ablative study on design choices w.r.t. joint selection, temporal edges, and constraints. The bottom part compares our final result with two strong baselines described in the text. HT:Head Top, N:Neck, S:Shoulders, W:Wrists, A:Ankles

crease in performance. This is most likely due to the weaker DeepMatching correspondences between hip joints, which are difficult to match. When only the body extremities (HT:W:A) are used for temporal edges, we obtain a similar MOTA as for (HT:N:S), but slightly worse other tracking measures. Considering the MOTA performance and the complexity of our graph structure, we use (HT:N:S) as our default setting.

Instead of considering only neighboring frames for temporal edges, we also evaluate the tracking performance while introducing longer-range temporal edges of up to 3 and 5 frames. Adding temporal edges between detections that are at most three frames ($\tau = 3$) apart improves the performance only slightly, whereas increasing the distance even further ($\tau = 5$) worsens the performance. For the rest of our experiments we therefore set $\tau = 3$.

To evaluate the proposed optimization objective (3) for joint multi-person pose estimation and tracking in more detail, we have quantified the impact of various kinds of constraints (10)-(15) enforced during the optimization. To this end, we remove one type of constraints at a time and solve the optimization problem. As shown in Tab. 2, all types of constraints are important to achieve best performance, with the spatial transitivity constraints playing the most crucial role. This is expected since these constraints ensure that we obtain valid poses without multiple joint types assigned to one person. Temporal transitivity and spatio-temporal con-

Method	Head	Sho	Elb	Wri	Hip	Knee	Ank	mAP
<i>Impact of the temporal connection density</i>								
HT	52.5	47.0	37.6	28.2	19.7	27.8	27.4	34.3
HT:N:S	56.1	51.3	42.1	31.2	22.0	31.6	31.3	37.9
HT:N:S:H	56.3	51.5	42.2	31.4	21.7	31.6	32.0	38.1
HT:W:A	56.0	51.2	42.2	31.6	21.6	31.2	31.7	37.9
<i>Impact of the length of temporal connection (τ)</i>								
HT:N:S ($\tau = 1$)	56.1	51.3	42.1	31.2	22.0	31.6	31.3	37.9
HT:N:S ($\tau = 3$)	56.5	51.6	42.3	31.4	22.0	31.9	31.6	38.2
HT:N:S ($\tau = 5$)	56.2	51.3	41.8	31.1	22.0	31.4	31.5	37.9
<i>Impact of the constraints</i>								
All	56.5	51.6	42.3	31.4	22.0	31.9	31.6	38.2
All \ spat. transitivity	7.8	10.1	7.2	4.6	2.7	4.9	5.9	6.2
All \ temp. transitivity	50.5	46.8	37.5	27.6	20.3	30.1	28.7	34.5
All \ spatio-temporal	42.3	40.8	32.8	24.3	17.0	25.3	22.4	29.3
<i>Comparison with the state-of-the-art</i>								
Ours	56.5	51.6	42.3	31.4	22.0	31.9	31.6	38.2
BBox-Detection [34]								
+ LJPA [17]	50.5	49.3	38.3	33.0	21.7	29.6	29.2	35.9
+ CPM [40]	48.8	47.5	35.8	29.2	20.7	27.1	22.4	33.1
DeeperCut [16]	56.2	52.4	40.1	30.0	22.8	30.5	30.8	37.5

Table 3: Quantitative evaluation of multi-person pose estimation (mAP). HT:Head Top, N:Neck, S:Shoulders, W:Wrists, A:Ankles

straints also turn out to be important to obtain good results. Removing either of the two significantly decreases the recall, resulting in a drop in MOTA.

Since we are the first to report results on the Multi-Person PoseTrack dataset, we also develop two baseline methods by using the existing approaches. For this, we rely on a state-of-the-art method for multi-person pose estimation in images [17]. The approach uses a person detector [34] to first obtain person bounding box hypotheses, and then estimates the pose for each person independently. We extend it to videos as follows. We first generate person bounding boxes for all frames in the video using a state-of-the-art person detector (Faster R-CNN [34]), and perform person tracking using a state-of-the-art person tracker [38] and train it on the training set of the Multi-Person PoseTrack Dataset. We also discard all tracks that are shorter than 7 frames. The final pose estimates are obtained by using the Local Joint-to-Person Association (LJPA) approach proposed by [17] for each person track. We also report results when Convolutional Pose Machines (CPM) [40] are used instead. Since CPM does not account for joint occlusion and truncation, the MOTA score is significantly lower than for LJPA. LJPA [17] improves the performance, but remains inferior w.r.t. most measures compared to our proposed method. In particular, our method achieves the highest MOTA and MOTP scores. The former is due to a significantly higher recall, while the latter is a result of a more precise part localization. Interestingly, the person bounding-box tracking based baselines achieve a lower number of ID switches. We believe that this is primarily due to the powerful multi-target tracking approach [38], which can handle person identities more robustly.




Figure 4: **Left** Impact of the the temporal edge density. **Middle** Impact of the length of temporal edges. **Right** Impact of different constraint types.

5.2. Frame-wise Multi-Person Pose Estimation

The results for frame-wise multi-person pose estimation (mAP) are summarized in Tab. 3. Similar to the evaluation for pose tracking, we evaluate the impact of spatio-temporal connection density, length of temporal connections and the influence of different constraint types. Having connections only between Head Top (HT) detections results in a mAP of 34.3%. As for pose tracking, introducing temporal connections for Neck and Shoulders (HT:N:S) results in a higher accuracy and improves the mAP from 34.3% to 37.9%. The mAP elevates slightly more when we also incorporate connections for hip joints (HT:N:S:H). This is in contrast to pose tracking where MOTA dropped slightly when we also use connections for hip joints. As before, inclusion of edges between all detections that are in the range of 3 frames improves the performance, while increasing the distance further ($\tau = 5$) starts to deteriorate the performance. A similar trend can also be seen for the impact of different types of constraints. The removal of spatial transitivity constraints results in a drastic decrease in pose estimation accuracy. Without temporal transitivity constraints or spatio-temporal constraints the pose estimation accuracy drops by more than 3% and 8%, respectively. This once again indicates that all types of constraints are essential to obtain better pose estimation and tracking performance.

We also compare the proposed method with the state-of-the-art approaches for multi-person pose estimation in images. Similar to [17], we use Faster R-CNN [34] as person detector, and use the provided codes for LJPA [17] and CPM [40] to process each bounding box detection independently. We can see that person bounding box based approaches significantly underperform as compared to the proposed method. We also compare with the state-of-the-art method DeeperCut [16]. The approach, however, requires the rough scale of the persons during testing. For this, we use the person detections obtained from [34] to compute the scale using the median scale of all detected persons.

Our approach achieves a better performance than all other methods. Moreover, all these approaches require an additional person detector either to get the bounding boxes

[17, 40], or the rough scale of the persons [16]. Our approach on the other hand does not require a separate person detector, and we perform joint detection across different scales, while also solving the person association problem across frames.

We also visualize how multi-person pose estimation accuracy (mAP) relates with the multi-person tracking accuracy (MOTA) in Fig. 4. Finally, Tab. 4 provides mean and median runtimes for constructing and solving the spatio-temporal graph along with the graph size for $k = 31$ frames over all test videos.

	Runtime (sec./frame)	# of nodes	# of spatial edges	# of temp. edges
Mean	14.7	2084	65535	12903
Median	4.2	1907	58164	8540

Table 4: Runtime and size of the spatio-temporal graph ($\tau = 3$, HT:N:S, $k = 31$), measured on a single threaded 3.3GHz CPU .

6. Conclusion

In this paper we have presented a novel approach to simultaneously perform multi-person pose estimation and tracking. We demonstrate that the problem can be formulated as a spatio-temporal graph which can be efficiently optimized using integer linear programming. We have also presented a challenging and diverse annotated dataset with a comprehensive evaluation protocol to analyze the algorithms for multi-person pose estimation and tracking. Following the evaluation protocol, the proposed method does not make any assumptions about the number, size, or location of the persons, and can perform pose estimation and tracking in completely unconstrained videos. Moreover, the method is able to perform pose estimation and tracking under severe occlusion and truncation. Experimental results on the proposed dataset demonstrate that our method outperforms other baseline methods.

Acknowledgments. The authors are thankful to Chau Minh Triet, Andreas Doering, and Zain Umer Javaid for the help with annotating the dataset. The work has been financially supported by the DFG project GA 1927/5-1 (DFG Research Unit FOR 2535 Anticipating Human Behavior) and the ERC Starting Grant ARCA (677650).

References

- [1] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014. 5, 6
- [2] M. Andriluka, S. Roth, and B. Schiele. People-tracking-by-detection and people-detection-by-tracking. In *CVPR*, 2008. 2
- [3] V. Belagiannis, S. Amin, M. Andriluka, B. Schiele, N. Navab, and S. Ilic. 3d pictorial structures revisited: Multiple human pose estimation. *TPAMI*, 2015. 2
- [4] K. Bernardin and R. Stiefelhagen. Evaluating multiple object tracking performance: The CLEAR MOT metrics. *Image and Video Processing*, 2008. 6
- [5] A. Bulat and G. Tzimiropoulos. Human pose estimation via convolutional part heatmap regression. In *ECCV*, 2016. 1, 2
- [6] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. In *CVPR*, 2016. 1, 2
- [7] J. Charles, T. Pfister, D. Magee, D. Hogg, and A. Zisserman. Personalizing human video pose estimation. In *CVPR*, 2016. 1, 5
- [8] X. Chen and A. L. Yuille. Parsing occluded people by flexible compositions. In *CVPR*, 2015. 1, 2
- [9] A. Cherian, J. Mairal, K. Alahari, and C. Schmid. Mixing Body-Part Sequences for Human Pose Estimation. In *CVPR*, 2014. 1, 2, 5
- [10] W. Choi. Near-online multi-target tracking with aggregated local flow descriptor. In *ICCV*, 2015. 6
- [11] M. Eichner and V. Ferrari. We are family: Joint pose estimation of multiple persons. In *ECCV*, 2010. 2, 5
- [12] G. Gkioxari, B. Hariharan, R. Girshick, and J. Malik. Using k-poselets for detecting people and localizing their keypoints. In *CVPR*, 2014. 1, 2
- [13] G. Gkioxari, A. Toshev, and N. Jaitly. Chained predictions using convolutional neural networks. In *ECCV*, 2016. 1, 2
- [14] P. Hu and D. Ramanan. Bottom-up and top-down reasoning with hierarchical rectified gaussians. In *CVPR*, 2016. 1, 2
- [15] E. Insafutdinov, M. Andriluka, L. Pishchulin, S. Tang, E. Levinkov, B. Andres, and B. Schiele. Articulated multi-person tracking in the wild. In *CVPR*, 2017. 2
- [16] E. Insafutdinov, L. Pishchulin, B. Andres, M. Andriluka, and B. Schiele. Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In *ECCV*, 2016. 1, 2, 4, 5, 7, 8
- [17] U. Iqbal and J. Gall. Multi-person pose estimation with local joint-to-person associations. In *ECCV Workshop on Crowd Understanding*, 2016. 1, 2, 4, 7, 8
- [18] U. Iqbal, M. Garbade, and J. Gall. Pose for action - action for pose. In *FG*, 2017. 1, 2
- [19] H. Izadinia, I. Saleemi, W. Li, and M. Shah. (MP)2T: Multiple people multiple parts tracker. In *ECCV*, 2012. 2
- [20] A. Jain, J. Tompson, Y. LeCun, and C. Bregler. Modeep: A deep learning framework using motion features for human pose estimation. In *ACCV*, 2014. 1, 2
- [21] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. Black. Towards understanding action recognition. In *ICCV*, 2013. 2, 5
- [22] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, 2010. 5
- [23] L. Ladicky, P. H. Torr, and A. Zisserman. Human pose estimation using a joint pixel-wise and part-wise formulation. In *CVPR*, 2013. 2
- [24] Y. Li, C. Huang, and R. Nevatia. Learning to associate: Hybridboosted multi-target tracker for crowded scene. In *CVPR*, 2009. 6
- [25] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [26] A. Milan, L. Leal-Taixe, I. Reid, S. Roth, and K. Schindler. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. 6
- [27] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016. 1, 2
- [28] D. Park and D. Ramanan. N-best maximal decoders for part models. In *ICCV*, 2011. 1, 2
- [29] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. In *ICCV*, 2015. 1, 2
- [30] L. Pishchulin, E. Insafutdinov, S. Tang, B. Andres, M. Andriluka, P. Gehler, and B. Schiele. DeepCut: Joint subset partition and labeling for multi person pose estimation. In *CVPR*, 2016. 1, 2, 4, 5, 6
- [31] L. Pishchulin, A. Jain, M. Andriluka, T. Thormählen, and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR*, 2012. 2
- [32] U. Rafi, I. Kostrikov, J. Gall, and B. Leibe. An efficient convolutional network for human pose estimation. In *BMVC*, 2016. 1, 2
- [33] V. Ramakrishna, T. Kanade, and Y. Sheikh. Tracking human pose by tracking symmetric parts. In *CVPR*, 2013. 1, 2
- [34] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 7, 8
- [35] B. Sapp, D. J. Weiss, and B. Taskar. Parsing human motion with stretchable models. In *CVPR*, 2011. 5
- [36] H. Shen, S.-I. Yu, Y. Yang, D. Meng, and A. Hauptmann. Unsupervised video adaptation for parsing human motion. In *ECCV*, 2014. 5
- [37] M. Sun and S. Savarese. Articulated part-based model for joint object detection and pose estimation. In *ICCV*, 2011. 2
- [38] S. Tang, B. Andres, M. Andriluka, and B. Schiele. Multi-person tracking by multicut and deep matching. In *ECCV Workshop on Benchmarking Multi-target Tracking*, 2016. 5, 7
- [39] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *CVPR*, 2014. 2
- [40] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016. 1, 2, 7, 8
- [41] P. Weinzaepfel, J. Revaud, Z. Harchaoui, and C. Schmid. DeepFlow: Large displacement optical flow with deep matching. In *ICCV*, 2013. 5
- [42] B. Yang and R. Nevatia. An online learned CRF model for multi-target tracking. In *CVPR*, 2012. 6

- [43] Y. Yang and D. Ramanan. Articulated human detection with flexible mixtures of parts. *TPAMI*, 2013. [2](#)
- [44] D. Zhang and Mubarak. Human pose estimation in videos. In *ICCV*, 2015. [1](#), [2](#)
- [45] W. Zhang, M. Zhu, and K. G. Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *ICCV*, 2013. [2](#), [5](#)
- [46] S. Zuffi, J. Romero, C. Schmid, and M. J. Black. Estimating human pose with flowing puppets. In *ICCV*, 2013. [1](#), [2](#)