JointFlow: Temporal Flow Fields for Multi Person Pose Tracking

Andreas Doering doering@iai.uni-bonn.de Umar Iqbal uiqbal@iai.uni-bonn.de Juergen Gall gall@iai.uni-bonn.de Computer Vision Group University of Bonn Bonn, DE

Abstract

In this work we propose an online multi person pose tracking approach which works on two consecutive frames I_{t-1} and I_t . The general formulation of our temporal network allows to rely on any multi person pose estimation approach as spatial network. From the spatial network we extract image features and pose features for both frames. These features serve as input for our temporal model that predicts Temporal Flow Fields (TFF). These TFF are vector fields which indicate the direction in which each body joint is going to move from frame I_{t-1} to frame I_t . This novel representation allows to formulate a similarity measure of detected joints. These similarities are used as binary potentials in a bipartite graph optimization problem in order to perform tracking of multiple poses. We show that these TFF can be learned by a relative small CNN network whilst achieving state-of-the-art multi person pose tracking results.

1 Introduction

Understanding of human body pose is an important information for many scene understanding problems such as activity recognition, surveillance and human-computer interaction. Estimating the pose in unconstrained environments with multiple interacting people is a challenging problem. Apart from the large amounts of appearance variation and complex human body articulation, it also poses additional challenges such as large scale variation within a single scene, varying number of persons and body part occlusion and truncation. Multiperson pose estimation in videos increases the complexity even further since it also requires to tackle the problems of person association over time, large person or camera motion, motion blur, etc.

In this work, we address the problem of multi-person pose tracking in videos, i.e, our goal is to estimate the pose of all persons appearing in the video and assign a unique identity to each person over time. The state-of-the-art approaches [12, 53] in this direction build on the recent progress in multi-person pose estimation in images and first estimate the poses from images using off-the-shelf methods followed by an additional step for person association over time. There exist two main approaches for person association. The *online* approach performs matching of the poses estimated at each time frame with the previously tracked

poses and assigns an identity to each pose before moving to the next time step. In contrast, offline or batch-processing based approaches [11] first estimate the poses in the entire video and then perform pose tracking while enforcing global temporal coherency of the tracks. In any case, both types of approaches require some metrics to measure the similarity between a pair of poses. The choice of the metrics and features used for matching plays a crucial role in the performance of these approaches. Recent methods for pose tracking [1] rely on nonparametric metrics such as head normalized Percentage of Correct Keypoints (PCKh) [and Object Keypoint Similarity (OKS) [2] between a pair of poses, Intersection over Union (IoU) between the bounding boxes tightly enclosing each body pose [12], similarity between the image features extracted from the person bounding boxes [12] or the optical flow information [17], 19, 13]. The location based metrics such as PCKh, OKS or IoU, on one hand, assume that the poses change smoothly over time, and therefore, struggle in case of large camera or body pose motion and scale variations due to camera zoom. On the other hand, appearance based similarity metrics or optical flow information cannot handle large appearance variations due to person occlusions or truncation, motion blur, etc. The offline approaches try to tackle these challenges by enforcing long-range temporal coherence. This is often done by formulating the problem using complex spatio-temporal graphs [12], [13] which results in very high inference time, and therefore, makes these methods infeasible for many applications.

In this work, we present an approach for online multi-person pose tracking. In contrast to existing methods that rely on task-agnostic similarity metrics, we propose a task-specific novel representation for person association over time. We refer to this representation as Temporal Flow Fields (TFF). TFF represent the movement of each body part between two consecutive frames using a set of 2D vectors encoded in an image. Our TFF representation is inspired by the Part Affinity Fields representation [**D**] that measures the spatial association between different body parts and is learned by a CNN. We integrate TFF in an online multiperson tracking approach and demonstrate that a greedy matching approach is sufficient to obtain state-of-the-art multi-person pose tracking results on the PoseTrack benchmark [**D**].

2 Related Work

Former incorporate person detectors and estimate the pose for each person proposal. For instance, Fang *et al.* [\blacksquare] extend a stacked hourglass network [\blacksquare] by two transformer networks, a spatial transformer network (STN) and a spatial de-transformer network (SDTN) respectively. Each person proposal is passed to the STN which automatically detects the person of interest and applies an affine transformation which centers the person in an upright position. After pose estimation the SDTN maps the pose back to the input image. In a final step, the best pose for each person proposal is selected by non-maxima suppression on poses. Chen *et al.* [\blacksquare] categorize invisible or occluded keypoints as "hard" whereas the remaining keypoints are classified as "simple". The proposed cascaded model reflects their

categorization of keypoints and is divided into two stages. The first stage is a feature pyramid network which detects "simple" keypoints (*e.g.* head). The second stage, called *RefineNet*, integrates all feature representations of different scales generated by the first stages. In that way, the *RefineNet* is able to incorporate enough context to detect occluded or invisible body parts.

Bottom-up approaches estimate the keypoints of all persons in a single run, but require a post-processing procedure to assemble these keypoints into person estimates. Cao *et al.* [**D**] estimate multiple poses in real-time. Their work extends the model proposed in [**D**] by introducing an additional branch which predicts vector fields between body parts of individual persons. These so called Part Affinity Fields (PAF) preserve location and orientation information [**D**] of limbs. The authors propose to use a greedy bipartite graph matching algorithm, which greedily connects joints that share the same body part. The work of Varadarajan *et al.* [**D**] introduces a more efficient greedy part assignment algorithm compared to [**D**]. After part belief maps and pairwise association maps are obtained like in [**D**], the number of part candidates is reduced to an approximate number of persons within a clustering step. By following the kinematic chain, body parts are assigned in a greedy fashion to joints of the most proximal candidate person cluster.

2.1 Multi-Person Pose Tracking

Even though a big advancement in multi person pose estimation in images has been achieved, very few works have addressed this problem in videos [2, 2, 2, 2, 2, 2], is one of the first works which tackels the problem of multi person pose estimation and tracking by solving a spatio-temporal graph matching problem. The spatio-temporal graph is created by densely connecting all detected joint candidates in the spatial domain. In the temporal domain all joints of the same class are connected. In order to find the best graph partition, a conditioned integer linear programming problem has to be optimized. For runtime reasons, [11] propose to sequentially optimize for temporal windows of a fixed size only. Nevertheless, the runtime is still too high which makes this work impractical for real-time applications. A very similar approach with comparable performance is proposed by [12] which in contrast to [12] relies on a sparse spatio-temporal graph. [12] propose a video pose estimation formulation which consists of a 3D extension of the Mask R-CNN model [1]. By integrating temporal information, the proposed model estimates person bounding boxes and poses which the authors refer to as *person tubes*. To achieve this, their network first predicts bounding boxes for each frame followed by a pre-trained Resnet-101 network [13] for pose estimation. In order to link the estimated poses in time, [12] propose to solve a bipartite graph matching problem in a greedy fashion and show that the achieved results are very close to the optimal solution obtained via Hungarian algorithm. By comparing different distance metrics, the authors show that Intersection over Union (IoU) of person bounding boxes achieves the best tradeoff between performance and runtime. Nevertheless, this approach requires to process entire sequences or portion of a sequence which limits the applicability for real-time applications.

In [\Box], the authors follow a very similar baseline as proposed in [\Box], but in contrast the authors rely on two different sources for person bounding boxes: a bounding box detector and optical flow. This allows to warp estimated poses of the previous frames $I_{\Delta t}$ with $\Delta t = \{1, \ldots, T\}$ into the current frame and a similarity metric between estimated and warped poses based on the Object Keypoint Similarity (OKS) is used for the calculation of binary potentials of a temporal graph. By utilizing greedy graph matching similar to [\Box] this approach achieves state-of-the-art results.



Figure 1: Proposed approach: For two consecutive input frames I_{t-1} and I_t , we utilize a Siamese network initialized by an arbitrary multi person pose estimation network. During spatial inference, pose features such as belief maps or part affinity fields are used to estimate the poses for each frame. Building on these pose features, our porposed temporal model predicts the temporal flow field for each detected joint which are used during inference to associate poses in time.

3 Overview

In this work, we propose to predict Temporal Flow Fields in an online fashion. To this extend, we evaluate two frames at a time as visualized in Figure 1 and estimate their poses. The structure of our temporal model allows to utilize any network architecture for the task of multi person pose estimation. In the context of this work, we use the CNN of $[\Box]$ as a component in our Siamese network. While the Siamese network is used to predict the poses in both frames, we take the last layer as input for the temporal CNN which predicts the Temporal Flow Fields (TFF). To track the poses, we then create a bipartite graph \mathcal{G} as illustrated in Figure 4b) from the estimated poses and use the estimated TFF as similarity measure (Sec. 4) in a bipartite graph matching problem.

4 Multi-Person Pose Tracking

We represent the body pose *P* of a person with *J* body joints as $P = \{p_j\}_{1:J}$, where $p_j = (x_j, y_j)$ represent the 2D pixel coordinates of the j^{th} body joint. Given an input video, our goal is to perform multi person pose estimation and tracking in an online manner. Formally, at every time instance *t* with video frame I_t containing N_t persons, we first estimate a set of poses $\mathcal{P}_t = \{P_t^1, \dots, P_t^{N_t}\}$ and then perform person association with the set of persons $\mathcal{P}_{t-1} = \{P_{t-1}^1, \dots, P_{t-1}^{N_t-1}\}$ tracked until the last video frame I_{t-1} . For pose estimation, we use an improved version of [**D**] that we will explain briefly in Sec. **5**. We formulate the problem of person association between the set of poses \mathcal{P}_t and \mathcal{P}_{t-1} as an energy maximization problem over a bipartite graph \mathcal{G} (Figure 4b) as follows

$$\hat{z} = \arg_{z} \sum_{P_{t} \in \mathcal{P}_{t}} \sum_{P_{t-1} \in \mathcal{P}_{t-1}} \Psi_{P_{t}, P_{t-1}'} \cdot z_{P_{t}, P_{t-1}'}$$
(1)

s.t.
$$\forall P_t \in \mathcal{P}_t, \sum_{P'_{t-1} \in \mathcal{P}_{t-1}} z_{P_t, P'_{t-1}} \le 1 \text{ and } \forall P'_{t-1} \in \mathcal{P}_{t-1}, \sum_{P_t \in \mathcal{P}} z_{P_t, P'_{t-1}} \le 1$$



Figure 2: Calculation of Temporal Flow Fields: Let $p_{j,k}^{t-1}$ and $p_{j,k}^{t}$ be the location of joint *j* of person *k* in frames I_{t-1} and I_t . For every point $p \in \Omega_{j,k}$ located on the flow field, the TFF $T_{j,k}^*(p)$ contains a unit vector v and 0 otherwise.

where $z_{P_t,P'_{t-1}} \in \{0,1\}$ is a binary variable which indicates that the poses $P_t \in \mathcal{P}_t$ and $P'_{t-1} \in \mathcal{P}_{t-1}$ are associated with each other, and the binary potentials $\Psi_{P_t,P'_{t-1}}$ define the similarity between the pair of poses P_t and P'_{t-1} .

4.1 Temporal Flow Fields

We model the binary potentials $\Psi_{P_t,P'_{t-1}}$ (1) by Temporal Flow Fields (TFF) and define each TFF as a vector field that contains a unit vector v for each pixel p = (x,y). Each unit vector $v = \frac{p'_{j,k} - p'_{j,k}}{\lambda_{j,k}}$ points towards the direction of the target joint location $p'_{j,k} \in P_t^k$ where $\lambda_{j,k} = ||p'_{j,k} - p'_{j,k}||_2$ is the Euclidean distance between the estimated joint locations of person k in frames I_{t-1} and I_t . We restrict the TFF to pixels that are close to the joint motion by a parameter σ and describe the set of pixels of the TFF as

$$\Omega_{j,k} = \{ p \mid 0 \le v \cdot (p - p_{j,k}^{t-1}) \le \lambda_{j,k} \land |v_{\perp} \cdot (p - p_{j,k}^{t-1})| \le \sigma \},$$
(2)

where v_{\perp} is a unit vector perpendicular to v as illustrated in Figure 2. This allows a pixel-wise definition of TFF for joint class j of person k

$$T_{j,k}^{*}(p) = \begin{cases} v & \text{if } p \in \Omega_{j,k} \\ 0 & \text{otherwise.} \end{cases}$$
(3)

In a final step, a single representation of a flow field T_j is generated for each joint class by aggregating the TFF among all estimated persons.

$$T_{j}^{*}(p) = \frac{1}{n_{t}(p)} \sum_{k=1}^{K} T_{j,k}^{*}(p),$$
(4)

where $n_t(p)$ is the number of non-zero unit vectors v at location p across all K persons.

4.1.1 Model

For the prediction of Temporal Flow Fields, we propose an efficient CNN as illustrated in Figure 3b) which consists of five 7×7 -convolution layers with a stride of one pixel followed by two 1×1 -convolutions. Non-linearity is achieved by ReLU layers after each convolution.



Figure 3: Proposed Model Structure: a) Siamese network to extract pose features (*SVGG*, *Belief*, *PAF*) for frames I_{t-1} and I_t and b) temporal network to extract temporal part affinity fields for feature map input SVGG and Beliefmaps (*SVGG* + *Belief*) from the two frames.

As input, the network expects image features and pose features. These are obtained from the Siamese network visualized in Figure 3a) which is initialized by a modified version of $[\square]$ consisting of six stages. In particular, image features of both frames I_{t-1} and I_t are obtained by a feature extraction layer as illustrated in Figure 3a). We refer to these image features as *SVGG*. Additionally, the Siamese network predicts beliefmaps and Part Affinity Fields (PAFs) (cf. $[\square]$) at each stage which we refer to as pose features. Based on image features and pose features, extracted from the last stage, the temporal model predicts TFF.

For the training of our model, we calculate the weighted squared L2 loss of the form

$$\mathcal{L} = \sum_{j=1}^{J} \sum_{p \in \Omega} M(p) \cdot \left\| T_{j}^{*}(p) - T_{j}(p) \right\|_{2}^{2},$$
(5)

where $T_j^*(p)$ and $T_j(p)$ are the ground truth TFF and the predicted TFF at pixel location p respectively. M is a binary mask with M(p) = 0 for all pixels located on an ignore region, i.e., a region for which the dataset does not provide any annotations.

4.2 Inference

During inference, we partition the bipartite graph \mathcal{G} by optimizing (1) using a greedy approach. In order to obtain the binary potentials Ψ_{P_{t-1},P'_t} we first generate J temporal bipartie subgraphs \mathcal{G}_j with a set of edges $Z_j^t = \left\{ z_j^{p_{t-1},p_t} \middle| P_t \in \mathcal{P}_t, P_{t-1} \in \mathcal{P}_{t-1} \right\}$ that connect all detected joints of class j in frame I_{t-1} with all detected joints in frame I_t of the same class. Figure 4a) illustrates such subgraphs. Along each temporal edge of \mathcal{G}_j , we follow the estimated TFF and obtain a flow field aggregate given by

$$E_{aggr}(p_{j,m}^{t-1}, p_{j,n}^{t}) = \int_{o=0}^{o=1} T_{j}(i(o))^{\top} \frac{(p_{j,n}^{t} - p_{j,m}^{t-1})}{\left\| p_{j,n}^{t} - p_{j,m}^{t-1} \right\|_{2}} do,$$
(6)

where $i(o) = (1-o) \cdot p_{j,m}^{t-1} + o \cdot p_{j,n}^{t}$ is a function that interpolates the location between both detected joints $p_{j,m}^{t-1}$ and $p_{j,n}^{t}$. The value is high if the TFF points in the same direction as



Figure 4: Temporal edge candidate generation for incomplete pose estimates: a) on a joint level and b) on a person level where $\Psi_{P_{t-1}^1, P_t^1}$ and $\Psi_{P_{t-1}^1, P_t^2}$ are the accumulated temporal edge potentials estimated according to Equation (8).

 $p_{j,n}^t - p_{j,m}^{t-1}$ along i(o). In addition to this formulation, we have to consider a special case: if there is no motion of a joint between frames, no flow field would exist and the flow field aggregate would be zero. To overcome this issue, we incorporate the Euclidean distance $\Delta p_{j,m,n}^t = ||p_{j,n}^t - p_{j,m}^{t-1}||_2$ between both joint locations into our similarity measure and define

$$E_{j}(p_{j,m}^{t-1}, p_{j,n}^{t}) = \begin{cases} E_{aggr}^{T}(p_{j,m}^{t-1}, p_{j,n}^{t}) & \text{if } \Delta p_{j,m,n}^{t} \ge \tau_{\Delta} \\ 1 & \text{if } \Delta p_{j,m,n}^{t} < \tau_{\Delta}, \end{cases}$$
(7)

where τ_{Δ} is a pre-defined distance threshold. This definition allows to formulate the binary potentials required to solve (1). However, instead of solving *J* different bipartite graph matching problems for each joint class *j*, we convert each estimated pose into a node of graph \mathcal{G} as illustrated in Figure 4b). The temporal potential $\Psi_{P_{t-1}^m,P_t^n}$ is then defined as the accumulated similarity between all joints *j* of persons P_{t-1}^m and P_t^n

$$\Psi_{P_{t-1}^m, P_t^n} = \sum_{j=1}^J \mathbb{1}(p_{j,m}^{t-1}, p_{j,n}^t) \cdot E_j(p_{j,m}^{t-1}, p_{j,n}^t),$$
(8)

where $\mathbb{1}(p_{j,m}^{t-1}, p_{j,n}^{t})$ is a binary function with $\mathbb{1}(p_{j,m}^{t-1}, p_{j,n}^{t}) = 1$ if both joints are detected. An example is shown in Figure 4a) and 4b): temporal edges between estimated joints of person $P_{t-1}^1 \in \mathcal{P}_{t-1}$ of frame t-1 and persons $P_t^1, P_t^2 \in \mathcal{P}_t$ in frame t are estimated. In this particular case, persons P_{t-1}^1 and P_t^1 share 12 temporal connections among joints whereas persons P_{t-1}^1 and P_t^2 share 9 temporal edges. The costs of these edges are accumulated and assigned as potential $\Psi_{P_{t-1}^1, P_t^1}$ and $\Psi_{P_{t-1}^1, P_t^2}$ to the temporal connections among persons as shown in Figure 4b).

By solving (1), we assign each detected person to either one of the poses from the previous frame, which continues the track, or a new track is initialized if no assignment is possible.

5 Implementation Details

For the spatial part of our proposed approach (Figure 1), we re-implemented the method of [**D**] and applied two minor modifications: instead of initializing the feature extraction part by 10 layers of VGG19, we increase the number of layers to 12. The second modification includes a different edge configuration for the prediction of Part Affinity Fields. Both changes

		MOTA	MOTP	Prec	Rec				MOTA	MOTP	Prec	Rec	mAP
Input Features	τ_{NMS}	Total	Total	Total	Total	Input Fe	eatures	τ_{NMS}	Total	Total	Total	Total	Total
SVGG SVGG + Belief SVGG + Belief + PAF	0.1 0.1 0.1	56.0 56.2 56.2	84.2 84.2 84.2	82.4 82.4 82.4	74.9 74.9 74.9	SVGG + SVGG + SVGG +	⊦ Belief + Belief ⊦ Belief	0.1 0.2 0.3	56.2 59.1 58.2	84.2 84.4 84.9	82.4 87.1 91.1	74.9 71.9 66.8	69.3 67.0 62.4

Table 1: Impact of different combinations of input Table 2: Impact of threshold τ_{NMS} during NMS offeatures on the pose tracking performance.beliefmaps on pose estimation and tracking perfor-

mance.

result in a gain in pose estimation performance. For further evaluations and a detailed description of the underlying edge configuration, we refer to the supplementary material. For the detection of joints, we perform Non-Maximum Suppression (NMS) on the estimated beliefmaps and discard all detections that do not meet a threshold $\tau_{NMS} = 0.2$. The spatial model was trained on the MSCOCO dataset [21] for 22 epochs with a learning rate of $\eta = 4 \cdot 10^{-5}$ and a decay in learning rate of $\gamma = 0.333$ after 7 epochs. For finetuning, we rely on the PoseTrack dataset [2] and train for 3000 iterations with a learning rate of $\eta = 10^{-5}$ and a decay in learning rate after 1000 iterations.

Our temporal model is trained on the PoseTrack dataset for 40 epochs with a learning rate $\eta = 4 \cdot 10^{-6}$ and a learning rate decay of $\gamma = 0.333$ after seven epochs. Additionally, we fix $\tau_{\Delta} = 2$ and select $\sigma = 1$ as the desired width of our Temporal Flow Fields.

During inference, we evaluate pairs of frames at four different scales (0.5, 1, 1.5 and 2) and average the estimated results. Spatial greedy bipartite graph matching is performed similar to [**b**] to estimate the poses first, followed by our proposed pose tracking approach. Since we did not focus on optimizing the runtime as performed in [**b**], inference requires on average 5.6 seconds for a pair of images on an I7-5820K @ 3.3 GHz and a single 1080TI. Unlike [**b**], we did not resize the videos of the PoseTrack dataset, which also contains HD quality frames, to a smaller resolution which leaves a lot of space for improvements in runtime.

6 Experiments

Within the first experiment, we tested different combinations of pose features (Table 1) in order to evaluate their performance. To this extend, each model was trained for 40 epochs on the PoseTrack dataset [2] and we rely on the metrics proposed in [22] in order to measure pose tracking performance. Certainly, spatial image features (SVGG) provide a strong cue for the prediction of temporal vector fields of each joint class. Additional knowledge is provided by the estimated belief maps (Belief) of both input frames. Part Affinity Fields (PAFs) [5] represent the skeletonal structure of the human body, and were expected to boost the performance even further. As Table 1 reveals, this is not the case. In order to reduce the number of input parameters, we use SVGG + Belief as desired input for the temporal model. In order to evaluate the impact of an increased receptive field, we explored the impact of multiple stages similar to the spatial model. Our experiments have shown that further stages do not have any impact on the final performance.

In an additional set of experiments, we evaluate different thresholds τ_{NMS} for non-maximum suppression of the heatmaps used for the detection of joint candidates. Even though a higher threshold results in less accurate pose estimates, more confident detection candidates result in stronger person tracks. According to Table 2, we select $\tau_{NMS} = 2$ as passable trade-off resulting in a boost in tracking performance.

	MOTA	MOTP	Prec	Rec
Baselines	Total	Total	Total	Total
PCKh	50.0	84.4	87.1	71.9
IoU	57.7	84.4	87.1	71.9
OKS	58.8	84.4	87.1	71.9
Optical Flow	58.5	84.4	87.1	71.9
Temporal Flow Fields	59.1	84.4	87.1	71.9
T 11 2 G				

Table 3: Comparison to different baselines

6.1 Comparison to Baselines

In an additional set of experiments, the performance of TFF is compared to different tracking metrics, namely Intersection over Union (IoU) of persons, PCKh [3], Object Keypoint Similarity (OKS) and optical flow based tracking. For this purpose, the temporal potential defined in (8) has to be adapted to

$$\Psi_{P_{t-1}^m, p_t^n} = \begin{cases} \text{IoU}(BB_{P_{t-1}^m}, BB_{P_t^n}) & \text{for IoU} \\ \text{PCKh}(P_{t-1}^m, P_t^n) & \text{for PCKh} \\ \text{OKS}(P_{t-1}^m, P_t^n) & \text{for OKS}, \end{cases}$$
(9)

where $BB_{P_{t-1}^m}$ and $BB_{P_t^n}$ are the bounding boxes for persons $P_{t-1}^m \in \mathcal{P}_{t-1}$ and $P_t^n \in \mathcal{P}_t$ in frames I_{t-1} and I_t respectively. The bounding boxes for each person are estimated from the detected poses. Table 3 summarizes the results. All three metrics can not compete with the proposed TFF.

Optical flow based tracking requires a different set of changes. First of all, we rely on the approach of [\square] in order to estimate the optical flow $f \in \mathcal{R}^{w \times h \times 2}$. Similar to Temporal Flow Fields, the optical flow is a vector field which can be used to predict the movement of each joint from frame I_{t-1} to frame I_t . In order to incorporate the optical flow into the greedy bipartite graph matching algorithm, the flow field aggregation energy (6) has to be adapted as follows:

$$E_{flow}^{T}(p_{j,m}^{t-1}, p_{j,n}^{t}) = e^{-\frac{\left\|p_{j,n}^{t} - (p_{j,m}^{t-1} + f(p_{j,m}^{t-1}))\right\|^{2}}{\sigma_{flow}^{2}}},$$
(10)

where σ_{flow} controls the tolerance radius to mistakes. In that way, optical flow vectors f which vote for locations close to $p_{j,n}^t$ still contribute significantly to the energy E_{flow}^T . Experiments have shown, that $\sigma_{flow} = 30$ performs best. Although the network for optical flow [13] is much larger and more expensive than our network for TFF, TFF outperform the optical flow.

6.2 Comparison to State-of-the-Art

For a comparison to the state-of-the-art, we compare to the results on the PoseTrack validation set reported in [12, 13], 13] and to the results on the PoseTrack test set taken from the PoseTrack challenge leaderboard [10]. On the validation set, we achieve a total MOTA of 59.1 which can be improved up to 59.8 after pruning tracks of a length smaller than 7 frames. We submitted our results to the official validation server and achieved the second place on the leaderboard with a final MOTA of 53.1.

		MOTA	Prec	Rec	mAP
Approach	Evaluation Set	Total	Total	Total	Total
FlowTrack [12]	val	65.4	85.5	80.3	76.7
TFF	val	59.1	87.1	71.9	69.3
TFF + pruning	val	59.8	87.8	71.1	66.7
PoseFlow [val	58.3	87.0	70.3	66.5
ProTracker [🗖]	val	55.2	88.1	66.5	60.6
FlowTrack [test	57.8	79.4	80.3	74.6
TFF + pruning	test	53.1	82.6	69.7	63.3
HMPT*	test	51.9	-	-	63.7
ProTracker [🗖]	test	51.8	-	-	59.6
PoseFlow [🛂]	test	51.0	78.9	71.2	63.0
MVIG*	test	50.8	-	-	63.2
BUTD2*	test	50.6	-	-	59.2
Trackend*	test	49.7	-	-	57.5
PoseTrack [2]	test	48.4	-	-	59.4
MIPAL*	test	46.3	-	-	69.9
SOPT-PT *	test	42.0	-	-	58.2
ML_Lab*	test	41.8	-	-	70.3
ICG*	test	32.0	-	-	51.2
IC_IBUG*	test	-190.1	-	-	47.6

Table 4: Comparison to state-of-the-art. Approaches marked with * have not been published yet.

7 Conclusions

In this work, we proposed a convolutional neural network architecture for the task of online multi person pose tracking. Our approach consists of two sub-networks: a spatial network for multi person pose estimation and a temporal network which predicts Temporal Flow Fields. TFF are used by a greedy temporal bipartite graph matching algorithm which associates estimated poses in two consecutive frames I_{t-1} and I_t . The results showed that a strong structural knowledge in form of image features and belief maps of both frames are crucial for a good performance of our temporal model. By relying on such feature input, our approach achieves state-of-the-art pose tracking results, even with a small network architecture. For this reason, in future work we will investigate stronger network architectures in order to produce stronger Temporal Flow Fields which are able to cope with additional challenges like occlusions and long-term dependencies.

8 Acknowledgments

The work has been financially supported by the DFG projects GA 1927/5-1 (DFG Research Unit FOR 2535 Anticipating Human Behavior) and the ERC Starting Grant ARCA (677650).

References

- PoseTrack Challenge Leaderboard. https://posetrack.net/leaderboard. php, 2018. [Online; accessed 13-July-2018].
- M. Andriluka, U. Iqbal, E. Ensafutdinov, L. Pishchulin, A. Milan, J. Gall, and Schiele B. PoseTrack: A benchmark for human pose estimation and tracking. In *CVPR*, 2018.

- [3] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [4] Vasileios Belagiannis, Sikandar Amin, Mykhaylo Andriluka, Bernt Schiele, Nassir Navab, and Slobodan Ilic. 3d pictorial structures revisited: Multiple human pose estimation. *TPAMI*, 38(10):1929–1942, October 2016.
- [5] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017.
- [6] Xianjie Chen and Alan L. Yuille. Parsing occluded people by flexible compositions. In *CVPR*, 2015.
- [7] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. *CVPR*, 2018.
- [8] A. Criminisi and J. Shotton. *Decision Forests for Computer Vision and Medical Image Analysis.* Springer Publishing Company, 2013.
- [9] Marcin Eichner and Vittorio Ferrari. We are family: Joint pose estimation of multiple persons. In ECCV, 2010.
- [10] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017.
- [11] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61(1):55–79, January 2005.
- [12] Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, and Du Tran. Detect-and-track: Efficient pose estimation in videos. *CVPR*, 2018.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In CVPR, 2016.
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *ICCV*, 2017.
- [15] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017.
- [16] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schiele. Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In ECCV, 2016.
- [17] Eldar Insafutdinov, Mykhaylo Andriluka, Leonid Pishchulin, Siyu Tang, Evgeny Levinkov, Bjoern Andres, and Bernt Schiele. ArtTrack: Articulated Multi-person Tracking in the Wild. In CVPR, 2017.
- [18] Umar Iqbal and Juergen Gall. Multi-person pose estimation with local joint-to-person associations. In *ECCV*, 2016.
- [19] Umar Iqbal, Anton Milan, and Juergen Gall. Posetrack: Joint multi-person pose estimation and tracking. In CVPR, 2017.

- [20] Lubor Ladicky, Philip H. S. Torr, and Andrew Zisserman. Human pose estimation using a joint pixel-wise and part-wise formulation. In *CVPR*, 2013.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [22] Anton Milan, Laura Leal-Taixé, Ian D. Reid, Stefan Roth, and Konrad Schindler. MOT16: A benchmark for multi-object tracking. *ArXiv-Preprint*, 2016.
- [23] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [24] Xuecheng Nie, Jiashi Feng, Junliang Xing, and Shuicheng Yan. Generative partition networks for multi-person pose estimation. *ECCV*, 2018.
- [25] George Papandreou, Tyler Zhu, Nori Kanazawa, Alexander Toshev, Jonathan Tompson, Chris Bregler, and Kevin Murphy. Towards accurate multi-person pose estimation in the wild. In *CVPR*, 2017.
- [26] Leonid Pishchulin, Arjun Jain, Mykhaylo Andriluka, Thorsten Thormaehlen, and Bernt Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR*, 2012.
- [27] Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Björn Andres, Mykhaylo Andriluka, Peter Gehler, and Bernt Schiele. Deepcut: Joint subset partition and labeling for multi person pose estimation. In *CVPR*, 2016.
- [28] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. LCR-Net: Localization-Classification-Regression for Human Pose. In *CVPR*, 2017.
- [29] Srenivas Varadarajan, Parual Datta, and Omesh Tickoo. A greedy part assignment algorithm for real-time multi-person 2d pose estimation. *ArXiv-Preprint*, 2017.
- [30] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *CVPR*, 2016.
- [31] Fangting Xia, Peng Wang, Xianjie Chen, and Alan L. Yuille. Joint multi-person pose estimation and semantic part segmentation. In *CVPR*, 2017.
- [32] B. Xiao, H. Wu, and Y. Wei. Simple Baselines for Human Pose Estimation and Tracking. ECCV, 2018.
- [33] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. Pose Flow: Efficient online pose tracking. *BMVC*, 2018.
- [34] Xiangyu Zhu, Yingying Jiang, and Zhenbo Luo. Multi-person pose estimation for posetrack with enhanced part affinity fields. Technical report, Samsung Research Beijing, 2017.

A Supplementary Material

A.1 Qualitative Results



Figure 5: Qualitative results for sequences of the PoseTrack validation set [].

A.2 Baseline Improvement



Figure 6: Different edge configurations used for the training of different spatial models.

We evaluate the robustness of different edge configurations as shown in Figure 6. This is motivated by the fact that edge configuration a) is prone to errors. If a single edge is not estimated correctly, the entire pose breaks. Similar to [\Box] we introduce skip connections to the standard model (Figure 6 b) *Bypass model*). Figure 6 c) illustrates a different idea to connect joints which we refer to as *Range of Motion (ROM) model* since pairs of joints are connected if both lie within the same ROM of a third joint. Further we train an edge configuration as proposed in [\Box] which we refer to as *Extended model*. For completeness, we introduce a *nearly-fully-connected (NFC) model* (Figure 6 e)) which connects most nearby

DOERING, IQBAL, GALL: JOINTFLOW

Model	VGG Layers	Trained on	Head	Shou	Elb	Wri	Hip	Knee	Ankl	Total mAP
Standard	12	MSCOCO + PoseTrack	82.9	80.3	69.9	59.0	67.8	59.2	51.4	68.3
Bypass	12	MSCOCO + PoseTrack	83.0	79.2	67.6	59.0	66.2	61.2	53.6	68.2
ROM	12	MSCOCO + PoseTrack	82.0	76.2	70.3	57.9	69.3	61.7	54.1	68.3
Extended	12	MSCOCO + PoseTrack	80.0	80.8	71.3	57.8	72.5	63.3	53.9	69.3
NFC	12	MSCOCO + PoseTrack	78.3	75.8	68.3	56.9	69.2	62.1	53.5	67.1

Table 5: The evaluation of different edge configurations reveals that the *Extended* edge configuration performs best compared to the *Standard* edge configuration.

joints. We rely on the metric proposed in [22] for the estimation of mean average precision (mAP) of all our pose estimation models. Table 5 shows the results, using $\tau_{NSM} = 0.1$. In all other experiments, we use the Extended model.