

# Efficient Online Multi-Person 2D Pose Tracking with Recurrent Spatio-Temporal Affinity Fields

Yaadhav Raaj    Haroon Idrees    Gines Hidalgo    Yaser Sheikh  
The Robotics Institute, Carnegie Mellon University  
{raaj@cmu.edu, hidrees@cs.cmu.edu, gines@cmu.edu, yaser@cs.cmu.edu}

## Abstract

We present an online approach to efficiently and simultaneously detect and track 2D poses of multiple people in a video sequence. We build upon Part Affinity Fields (PAF) representation designed for static images, and propose an architecture that can encode and predict Spatio-Temporal Affinity Fields (STAF) across a video sequence. In particular, we propose a novel temporal topology cross-linked across limbs which can consistently handle body motions of a wide range of magnitudes. Additionally, we make the overall approach recurrent in nature, where the network ingests STAF heatmaps from previous frames and estimates those for the current frame. Our approach uses only online inference and tracking, and is currently the fastest and the most accurate bottom-up approach that is runtime-invariant to the number of people in the scene and accuracy-invariant to input frame rate of camera. Running at  $\sim 30$  fps on a single GPU at single scale, it achieves highly competitive results on the PoseTrack benchmarks.<sup>1</sup>

## 1. Introduction

Multi-person human pose estimation has received considerable attention in the past few years assisted by deep convolutional learning as well as COCO [21] and MPII [3] datasets. The recently introduced PoseTrack dataset [17] has provided the community with a large scale corpus of video data with multiple people in the scenes. In this paper, our aim is to utilize these towards building a truly online and real-time multi-person 2D pose estimator and tracker that is deployable and scalable while achieving high performance and requiring minimal post-processing. The potential uses include real-time and closed-loop applications with low latency where the execution is in sync with frame rate of camera such as self-driving cars and augmented reality.

The real-time and online nature of such an approach introduces several challenges: i) scenes with multiple peo-

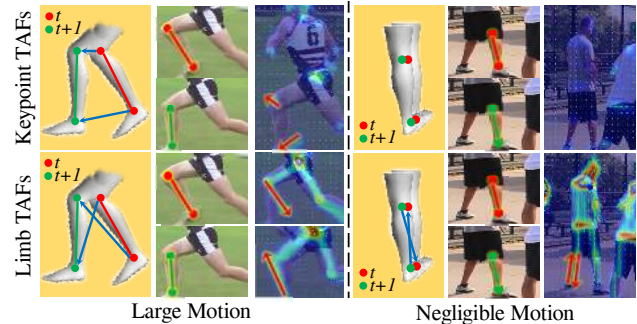


Figure 1: We solve multi-person human pose tracking by encoding change in position and orientation of keypoints or limbs across time as Temporal Affinity Fields (TAFs) in a recurrent fashion. **Top:** Modeling TAFs (blue arrows) through keypoints works when motion occurs but fails during limited motion making temporal association difficult. **Bottom:** Cross-linked TAFs across limbs perform consistently for all kinds of motions providing redundancy and smoother encoding for further refinement and prediction.

ple demand handling of occlusion, proximity and contact as well as limb articulation, and ii) it should be runtime-invariant to the number of people in the scene. Furthermore, iii) it must be capable of handling challenges induced from video data, such as large camera motion and motion blur across frames. We build upon the Part Affinity Fields (PAFs) [6] to overcome these challenges, which represent connections across body keypoints in static images as normalized 2D vector fields with position and orientation. In this work, we propose Temporal Affinity Fields (TAFs) which encode connections between keypoints across frames, including a unique cross-linked limb topology as seen in bottom row of Figure 1. In the absence of motion or when there is not enough data from previous frames, TAFs constructed between same keypoints, e.g., wrist-wrist or elbow-elbow across frames lose all associative properties (see top row of Fig. 1). In this case, the nullification of magnitude and orientation provides no useful information to dis-

<sup>1</sup>Project Page

cern between the case where a new person appears or where an existing person stops moving. This effect is compounded if these two cases occur in proximity together. However, the longer limb TAF connections allow information preservation even in the absence of motion or appearance of new people by preventing corruption of valid information with noise as the magnitude of motion becomes small. In the limiting case of zero motion, *the TAF effectively collapses to a PAF*. From the perspective of a network, TAF between keypoints destroys spatial information about keypoints as motion ceases, whereas TAF across keypoints simply learns to propagate the PAF, which is a much simpler task.

Furthermore, we work on videos in a recurrent manner to make the approach real-time, where computation of each frame leverages information from previous frames thereby reducing overall computation. Where the single-image pose estimation methods use multiple stages to refine heatmaps [6, 24], we exploit the redundant information in the video frames and divert the resources towards efficient computation of both poses and tracks across multiple frames. Thus, the multi-stage computation over images is divided over multiple frames in a video. Overall, we call this Recurrent Spatio-Temporal Affinity Fields (STAF) and it achieves highly competitive results on the PoseTrack benchmarks: [64.6% mAP, 58.4% MOTA] on single scale at  $\sim 30$  FPS, and [71.5% mAP, 61.3% MOTA] on multiple scales at  $\sim 7$  FPS on the PoseTrack 2017 validation set using one GTX 1080 Ti. As of writing, our approach currently ranks second for accuracy and at third place for tracking on the 2017 challenge [1]. Note that, our tracking approach is truly online on a per-frame basis with no post processing.

The rest of the paper is organized as follows. In Sec. 2, we discuss related work and situate the paper in the literature. In Sec. 3, we present details of our approach, training procedure as well as tracking and inference algorithm. Finally, we present results and ablation experiments in Sec. 4 and conclude the paper in Sec. 5.

## 2. Related Work

Early methods for human pose estimation localized keypoints or body parts of individuals but did not consider multiple people simultaneously [4, 28, 36, 20, 33]. Hence, these methods were not adept at localizing keypoints of highly articulated or interacting people. Person detection was typically used which followed single-person keypoint detection [29, 11, 32, 16]. With deep learning, human detection methods such as Mask-RCNN [10, 14] were employed to directly predict multiple human bounding boxes through ROI-pooling followed by pose estimation per person [12]. However, these methods suffered when people were in close proximity as bounding boxes got grouped together. Furthermore, these **top-down** methods required more computation as the number of people increased in the image, making

them inadequate for real-time pose estimation and tracking.

The **bottom-up** Part Affinity Fields (PAF) method [6] produced a spatial encoding of pair-wise body part connections in the image space, followed by greedy bipartite graph matching for inference permitting consistent computation speed irrespective of the number of people. Person Lab [25] built upon these ideas to incorporate redundant connections on people with a less greedy inference approach getting highly competitive results on the COCO [22] and MPII [3] datasets. These methods work on single images and do not incorporate any keypoint tracking or past information.

Many offline methods have been proposed to enforce temporal consistency of poses in videos [15, 17, 34]. These require solving spatio-temporal graphs or incorporating data from future frames making them inadequate for on-line operation. Alternatively, Song *et al.* and Pfister *et al.* [27, 31] demonstrate how optical flow fields could be predicted per keypoint by formulating the input to be multi-framed. LSTM Pose Machines [23] built upon previous work demonstrating use of single stage per frame for video sequences. However, these networks did not model spatial relationship between keypoints and were evaluated on the single person Penn Action [37] and JHMDB [18] datasets.

A different line of works explored maintaining temporal graphs in neural networks for handling multiple people [9, 8]. Rohit *et al.* demonstrated that a 3D extension of Mask-RCNN, called person tubes, can connect people across time. However, this required applying grouped convolutions over a stack of frames reducing speed, and did not achieve better results for tracking than the Hungarian Algorithm baseline. Joint Flow [8] used the concept of Temporal Flow Field which connected keypoints across two frames. However, it did not use a recurrent structure and explicitly required a pair of images as input increasing run-time significantly. The flow representation also suffered from ambiguity when subjects moved slowly or were stationary and required special handling of such cases during tracking.

Top-down pose and tracking methods [34, 33, 7, 26, 14] have dominated the detection and tracking tasks [34] [35] in PoseTrack but their speed suffered due to explicit human detection and follow-up keypoint detection for each person. Moreover, modeling long-term spatio-temporal graphs for tracking in an offline manner hurts real-time applications. None of these methods are able to report any significant runtime-to-performance measures as they cannot run in real time. In this work, we demonstrate this problem can be solved in a simple elegant single-stage network that incorporates recurrence by using the previous pose heatmaps to predict both keypoints and their spatio-temporal associations. We call this Recurrent Spatio-Temporal Affinity Fields (STAF) which not only represents the prediction of Spatial (PAF) and Temporal (TAF) Affinity Fields but also how they are refined through past information.

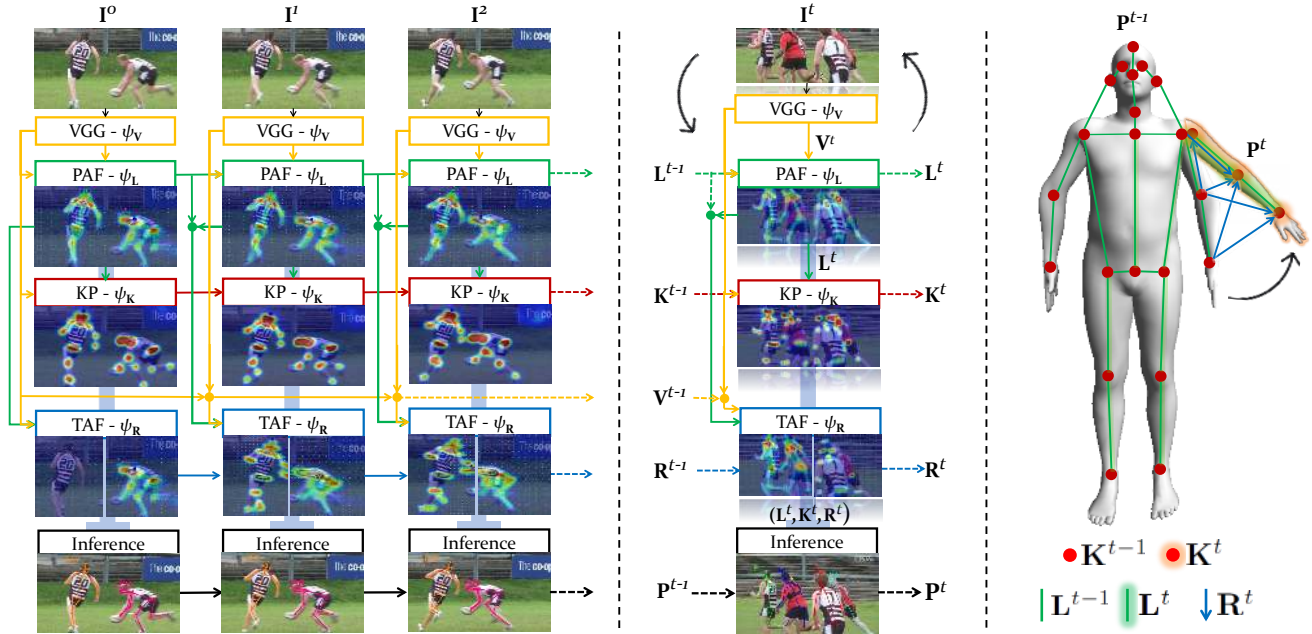


Figure 2: **Left:** Training architecture for one of our models which ingests video sequences in a recurrent manner across time while generating keypoints and connections across keypoints in each frame as Part Affinity Fields (PAFs), and connections between keypoints across frames as Temporal Affinity Fields (TAFs). Together, we call this Recurrent Spatio-Temporal Affinity Fields (STAF). Each module ingests outputs from other modules in both previous and current frames (shown with arrows) and refines it. **Center:** During inference, our network operates on a single video frame at each time step using past information. **Right:** During inference, we use the predicted heatmaps to detect and track people. Keypoints (red) are extracted first, then associated into poses and tracklets using PAFs (green), TAFs (blue), and tracklets from previous frames.

### 3. Proposed Approach

Our approach aims to solve the problems of keypoint estimation and tracking simultaneously in videos. We employ Recurrent Convolutional Neural Networks which we construct from four essential building blocks. Let  $P^t$  represent the pose of a person in a particular frame or time  $t$ , consisting of keypoints  $K = \{K_1, K_2, \dots, K_K\}$ . The Part Affinity Fields (PAFs)  $L = \{L_1, L_2, \dots, L_L\}$  are synthesized from keypoints in each frame. For tracking keypoints across frames a video, we propose Temporal Affinity Fields (TAFs) given by  $R = \{R_1, R_2, \dots, R_R\}$  which capture the recurrence and connect the keypoints across frames. Together, they are referred to as Spatio-Temporal Affinity Fields (STAF). These blocks are visualized in Fig. 2 where each block is shown with a different color: the raw convolutional feature from VGG backbone [30] are shown in amber, PAFs in green, keypoints in red and TAFs in blue.

Thus, the output of VGG backbone, PAFs, keypoints and TAFs are given by  $V$ ,  $L$ ,  $K$  and  $R$ , respectively, and computed through CNNs by  $\psi_V$ ,  $\psi_L$ ,  $\psi_K$  and  $\psi_R$ , respectively. The keypoint heatmaps are constructed from ground truth by placing a Gaussian kernel at the location of the annotated keypoint, whereas the PAFs and TAFs are constructed from

ground truth between pairs of keypoints for each person:

$$\tilde{L}_{k \rightarrow k'}^t := \Omega(\tilde{K}_k^t, \tilde{K}_{k'}^t), \quad \tilde{R}_{k \rightarrow k'}^t := \Omega(\tilde{K}_k^{t-1}, \tilde{K}_{k'}^t), \quad (1)$$

where  $\sim$  denotes the ground truth and the function  $\Omega(\cdot)$  places a directional unit vector at every pixel within a pre-defined radius of the line connecting the two keypoints.

#### 3.1. Video Models for Pose Estimation and Tracking

Next, we present the three models comprising the four blocks capable of estimating keypoints and STAF. The input to each network consists of a set of consecutive frames of a video. Each block in each network consists of five  $7 \times 7$  and two  $1 \times 1$  convolution layers. Each  $7 \times 7$  layer is replaceable with the concatenation of three  $3 \times 3$  convolution layers providing the same receptive field. The first stage has a unique set of weights from subsequent frames as it cannot incorporate any previous data and also has a lower depth which was found to improve results (see Sec. 4). The VGG features are computed for each frame. For frame  $I^t$  at time  $t$  of the video, they are computed as  $V^t = \psi_V(I^t)$ .

**Model I:** Given  $V^{t-1}$  and  $V^t$ , the the following equations

describe the first model:

$$\begin{aligned} \mathbf{L}^t &= \psi_{\mathbf{L}}(\mathbf{V}^t, \psi_{\mathbf{L}}^{q-1}(\cdot)), \\ \mathbf{K}^t &= \psi_{\mathbf{K}}(\mathbf{V}^t, \psi_{\mathbf{L}}^q(\cdot), \psi_{\mathbf{K}}^{q-1}(\cdot)), \\ \mathbf{R}^t &= \psi_{\mathbf{R}}(\mathbf{V}^{t-1}, \mathbf{V}^t, \mathbf{L}^{t-1}, \mathbf{L}^t, \mathbf{R}^{t-1}), \end{aligned} \quad (2)$$

where  $\psi^q$  means  $q$  recursive applications of  $\psi$ . In our experiments, we found that performance plateaus at  $q = 5$ . In Model I, PAFs are obtained by recursive application of  $\psi_{\mathbf{L}}$  on concatenated input from VGG features and PAFs from previous stage. Similarly, keypoints depend on VGG features, keypoints from the previous stage and PAFs from the current stage. Finally, TAFs are dependent on VGG features and PAFs from both the previous and current frames, as well as TAFs from previous frame. This model produces good results but is the slowest due to recursive stages.

**Model II:** Unlike Model I with multiple applications of CNNs for PAFs and keypoints, Model II computes the PAFs and keypoints in a single pass as visualized in Fig. 2:

$$\begin{aligned} \mathbf{L}^t &= \psi_{\mathbf{L}}(\mathbf{V}^t, \mathbf{L}^{t-1}), \\ \mathbf{K}^t &= \psi_{\mathbf{K}}(\mathbf{V}^t, \mathbf{L}^t, \mathbf{K}^{t-1}), \\ \mathbf{R}^t &= \psi_{\mathbf{R}}(\mathbf{V}^{t-1}, \mathbf{V}^t, \mathbf{L}^{t-1}, \mathbf{L}^t, \mathbf{R}^{t-1}). \end{aligned} \quad (3)$$

Replacing five stages with a single stage is expected to drop performance. Therefore, the multi-stage computation of PAFs and keypoints in Model II is supplanted with output of PAFs and keypoints from the previous frames. This boosts up the speed significantly without major loss in performance as it takes advantage of the redundant information in videos, i.e., the PAFs and keypoints from previous frame are a reliable guide to the location of PAFs and keypoints in the current frame.

**Model III:** Finally, the third model attempts to estimate Part and Temporal Affinity Fields through a single CNN:

$$\begin{aligned} [\mathbf{L}, \mathbf{R}]^t &= \psi_{[\mathbf{L}, \mathbf{R}]}(\mathbf{V}^{t-1}, \mathbf{V}^t, [\mathbf{L}, \mathbf{R}]^{t-1}), \\ \mathbf{K}^t &= \psi_{\mathbf{K}}(\mathbf{V}^t, \mathbf{L}^t, \mathbf{K}^{t-1}), \end{aligned} \quad (4)$$

where  $[\mathbf{L}, \mathbf{R}]$  implies simultaneous computation of Part and Temporal Affinity Fields through a single CNN. For Model III, the channels corresponding to PAFs are then passed for keypoint estimation along with VGG features from current frame and keypoints from previous frame. As Model III consists of only three blocks, it has the fastest inference, however it proved to be the most difficult to train.

### 3.2. Topology of Spatio-Temporal Affinity Fields

For our body model, we define  $K = 21$  body parts or keypoints which is the union of body parts in COCO and MPII pose datasets. They include ears, nose and eyes from COCO; and head and neck from MPII. Next, there are several possible ways to associate and track the keypoints and




Figure 3: This figure illustrates the three possible topology variations for Spatio-Temporal Affinity Fields including the new cross-linked limb topology (b). Keypoints, PAFs and TAFs are represented by solid circles, straight lines and arrows, respectively.

STAF across frames as illustrated in Figure 3. In this figure, solid circles represent keypoints while straight lines and arrows stand for PAFs and TAFs, respectively. Figure 3(a) consists of TAFs between same keypoints as well as PAFs. For this topology, the number of TAFs and PAFs is 21 and 48, respectively. The TAFs capture temporal connections directly across keypoints similar to [8].

On the other hand, Figure 3(b) consists of TAFs between different limbs in a cross-linked manner across frames. The number of PAFs and TAFs is 48 and 96, respectively. We also tested the topology in Figure 3(c) which consists of 69 keypoints and limb TAFs only. This does not model any spatial links within frames across keypoints.

### 3.3. Model Training

During training, we unroll each model to handle multiple frames at once. Each model is first pre-trained in **Image Mode** where we present a single image or frame at each time instant to the model. This implies multiple applications of PAF and keypoint stages to the same frame. We train with COCO, MPII and PoseTrack datasets with a batch distribution of 0.7, 0.2 and 0.1, respectively, which corresponds to dataset sizes where each batch consists of images or frames from one dataset exclusively. For masking out un-annotated keypoints, we use the head bounding boxes available in MPII and PoseTrack datasets, and location of annotated keypoints for batches from COCO dataset. The net takes in  $368 \times 368$  images and has scaling, rotation and translation augmentations. Heatmaps are computed with an  $\ell_2$  loss with a stride of 8 resulting in  $46 \times 46$  dimensional heatmaps. We initialize the limb TAFs with PAFs in topology 3(b,c), and keypoint TAFs with zeros in topology 3(a,c). We train the net for a maximum of 400k iterations.

Next, we proceed training in the **Video Mode** where we expose the network to video sequences. For static image datasets including COCO and MPII, we augment data with video sequences that have length equal to number of times the network is unrolled by synthesizing motion with scaling, rotation and translation. We train COCO, MPII and PoseTrack in Video Mode with a batch distribution of 0.4, 0.1 and 0.5, respectively. Moreover, we also use skip-frame augmentation for video-based PoseTrack dataset where some of the randomly selected sequences skip up to 3 frames. We lock the weights of VGG module in Video Mode. For Model I, we only train the TAF module when training on videos. For Model II, we train keypoint, PAF and TAF modules for 5000 epochs, then lock all modules except TAF. In Model III, both STAF and keypoints remain unlocked throughout the 300k iterations.

### 3.4. Inference and Tracking

The method described till now predicts heatmaps of keypoints and STAF at every frame. Next, we present the framework to perform pose inference and tracking across frames given the predicted heatmaps. Let the inferred poses at time  $t$  be given by  $\{\mathbf{P}^{t,1}, \mathbf{P}^{t,2}, \dots, \mathbf{P}^{t,N}\}$  where the second superscript indexes over people at frame  $t$ . Each pose at a particular time consists of up to  $K$  keypoints that become part of a pose post inference, i.e.,  $\mathbf{P}^{t,n} = \{\bar{\mathbf{K}}_1^{t,n}, \bar{\mathbf{K}}_2^{t,n}, \dots, \bar{\mathbf{K}}_K^{t,n}\}$ .

The detection and tracking procedure begins with localization of keypoints at time  $t$ . The inferred keypoints  $\bar{\mathbf{K}}^t$  are obtained by rescaling the heatmaps to original image resolution followed by non-maximal suppression. Then, we infer PAF weights,  $\bar{\mathbf{L}}^t$ , and those for TAF,  $\bar{\mathbf{R}}^t$ , between all pairs of keypoints in each frame defined by the given topology, i.e.,

$$\bar{\mathbf{L}}_{k \rightarrow k'}^t = \omega(\bar{\mathbf{K}}_k^t, \bar{\mathbf{K}}_{k'}^t), \quad \bar{\mathbf{R}}_{k \rightarrow k'}^t = \omega(\bar{\mathbf{K}}_k^{t-1}, \bar{\mathbf{K}}_{k'}^t), \quad (5)$$

where the function  $\omega(\cdot)$  samples points between the two keypoints, computes the dot product between the the mean vector of the sampled points and the directional vector from the first to the second keypoint.

Both the inferred PAFs and TAFs are sorted by their scores before inferring the complete poses and associating them across frames with unique ids. We perform this in a bottom-up style where we utilize poses and inferred PAFs from the previous frame to determine the update, addition or deletion of tracklets. Going through each PAF in the sorted list, (i) we initialize a new pose if both keypoints in the PAF are unassigned, (ii) add to existing pose if one of the keypoints is assigned, (iii) update score of PAF in pose if both are assigned to the same pose, and (iv) merge two poses if keypoints belong to different poses with opposing keypoints unassigned. Finally, we assign id to each pose in the current




Figure 4: (a) Ambiguity when selecting between two wrist locations B and E is resolved by reweighing PAFs through TAFs. (b)-(d): With transitivity, incorrect PAFs containing ankles (c) are resolved with past pose (b) resulting in (d).

frame with the most frequent id of keypoints from the previous frame. For cases where we have ambiguous PAFs, i.e., multiple equally likely possibilities as seen in Figure 4, we use transitivity that reweighs PAFs with TAFs to disambiguate between them, using  $\alpha$  as a biasing weight. In this figure, keypoint  $\{A\}$  - an elbow - is under consideration with wrists  $\{B\}$  and  $\{E\}$  as two possibilities. We select the strongest TAF where  $\{A, B, C, D, A\}$  has a higher weight than  $\{A, E, F, G, A\}$ , computed as:

$$\bar{\mathbf{L}}_{k \rightarrow k'}^{t,n} = (1 - \alpha) * \omega(\bar{\mathbf{K}}_k^{t-1,n}, \bar{\mathbf{K}}_{k'}^{t,n}) + \alpha * \omega(\bar{\mathbf{K}}_k^{t,n}, \bar{\mathbf{K}}_{k'}^{t,n}).$$

## 4. Experiments

In this section, we present results of our experiments. Input images to networks are resized at  $W \times 368$  maintaining aspect ratio for single scale (SS); and  $W \times 736$ ,  $W \times 368$  and  $W \times 184$  for multiple scales (MS). The heatmaps for multiple scales are re-sized back to  $W \times 736$  and merged through averaging. This is followed by inference and tracking.

### 4.1. Ablation Study

We conducted a series of ablation studies to determine the construction of our network architecture:

**Filter Sizes:** As discussed in Sec. 3, each block either consists of five  $7 \times 7$  layers followed by two  $1 \times 1$  layers [6], or each  $7 \times 7$  layer is replaced with three  $3 \times 3$  layers similar to [5] in the alternate experiment. The results are shown in Table 1. We run single frame inference on Model I and find the  $3 \times 3$  filter size to be 2% more accurate than  $7 \times 7$ , with significant boosts in average precision of knee and ankle keypoints. It is also 40% faster while requiring 40% more memory.

**Video Mode / Depth of First Stage:** Next, we report results when training in *Image Mode (Im)* using single images, and when we continue training beyond images while exposing the network to videos and augmenting with synthetic motion in the *Video Mode (Vid)*. During testing, the network is run recurrently on video sequences with one frame per

Method	Hea	Sho	Elb	Wri	Hip	Kne	Ank	mAP	fps
Model I - 3x3	75.7	73.9	67.8	56.3	66.8	62.3	56.9	66.3	14
Model I - 7x7	76.0	73.3	66.4	54.0	63.4	59.2	52.2	64.3	10

Table 1: This table shows results for experiments with the two filter sizes on PoseTrack 2017 validation set.




Figure 5: Improvement in quality of heatmaps before (a,c) and after (b,d) the network is exposed to videos and synthetic motion augmentation. We observe better peaks and less noise across both PAF and keypoint heatmaps.

stage. Model II is deployed for these experiments. We find that by exposing the network to video sequences for 5000 iterations, we were able to boost the mAP as seen in Table 2 and Fig. 5. We also find that if we use the same depth, i.e., number of channels for the first frame as the other frames (128-128), the network was not able to generalize well to recurrent execution (56.6 mAP) when trained with Image Mode. When reducing the depth for the first frame to one-half, i.e. (64-128), we found that the generalization to videos was better (62.6 mAP). When trained with Video Mode, mAP increased further to 64.1. We reason that the 64-depth modules produced relatively vague outputs which gave sufficient room for the subsequent modules in the following frames to process and refine the heatmaps yielding a boost in performance. Furthermore, this also highlights the importance of incorporating shot change detection and running the first stage at each shot change.

Method	Hea	Sho	Elb	Wri	Hip	Kne	Ank	mAP	fps
Im - 7x7 - 128-128	74.6	69.6	55.5	40.2	56.4	47.2	44.0	56.6	27
Vid - 7x7 - 128-128	76.2	71.6	64.5	51.9	62.6	59.3	52.5	63.6	27
Im - 7x7 - 64-128	73.5	72.2	63.8	52.1	62.7	57.3	51.1	62.6	27
Vid - 7x7 - 64-128	75.8	73.4	65.5	53.8	64.2	58.4	51.4	64.1	27
Im - 3x3 - 64-128	73.5	72.5	65.0	52.7	63.7	57.7	53.2	63.4	35
Vid - 3x3 - 64-128	75.4	73.2	67.4	55.0	63.9	58.4	53.5	64.6	35

Table 2: This table shows single-scale performance using Model II before and after training with videos, filter sizes, as well as different depths for first stage.

**Effect of Camera Frame Rate on mAP:** For these experiments, we studied how the frame rate of the camera and number of stages affect the accuracy of pose estimation. With a high frame rate, the apparent motion between frames is smooth, which becomes relatively abrupt at low frame-rates. Therefore, the heatmaps from previous frames would

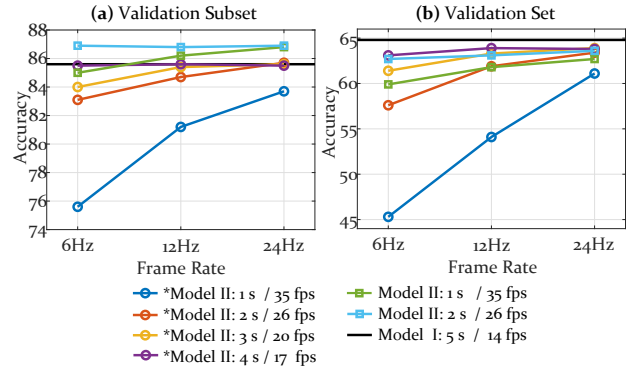


Figure 6: These graphs show mAP curves as a function of frame rates of camera, i.e., the rate at which an original 24Hz video is input to the method. The flat black line shows the performance of five-stage Model I, while “\*” in the legend indicates training using Image Mode only.

not be as useful at low frame-rates. We tested this hypothesis with Model I (five stages of the same modules without ingesting previous frame heatmaps), and Model II (different number of stages with each ingesting heatmaps from previous frame). We also evaluate the influence of training with Image and Video modes in Figure 6.

Fig. 6(a) shows results on a subset of ten sequences where the human subjects comprised at least 30% of the frame height in the PoseTrack 2017 validation set. Fig. 6(b) presents results on the entire validation set. The original videos were assumed to run at the film-standard 24 Hz, hence we ran experiments by varying frame rates at 24, 12 and 6 Hz through sub-sampling. The ground truth has been annotated at 6 Hz. As expected, accuracy is proportional to video frame rate and number of stages. When the Model II was trained in Image Mode, we observed small increments in accuracy until at four stages, it peaks at the same level as Model I. Upon training with Video Mode, it surpasses this accuracy peaking earlier at two stages.

When considering the entire validation set, the approach is still able to reap the benefits of more stages and training in Video Mode as can be seen in Fig. 6(b). However, it was barely able to reach the accuracy of the much slower Model I. For the validation set, the accuracy reduced when including sequences with smaller apparent size of humans. These sequences usually were more crowded as well, and passing in the previous heatmaps seemed to hurt the performance. The body parts of small-sized humans only occupied a few pixels in the heatmaps and the normalized direction vectors were inconsistent and random across frames.

**Influence of Topology / Model Type in Tracking:** Next, we report experiments on different combinations of topology defined in Fig. 3 with the three models presented in Sec. 3.1, both for pose estimation and tracking evaluated

using mean Average Precision (mAP) and Multiple Object Tracking Accuracy (MOTA) metrics in Table 3. We found an improvement in tracking using limb TAFs in Topology B versus keypoint TAFs in Topology A. As highlighted in Fig. 1, Topology A lacks associative properties when a keypoint has minimal motion or when a new person appears. Although we enforced spatial constraint that joint locations should be close in consecutive frames, and adjusted it according to scale (similar to [8]), this still resulted in false positives since it is difficult to disambiguate between a newly detected person and some nearby stationary person. Furthermore, where motion of a person tended to be small, Topology A resulted in jittery and noisy vectors causing more reliance on pixel distances. This was further exacerbated by recurrence where accumulation of noisy vectors from previous frame heatmaps deteriorated associative ability of Temporal Affinity Fields. Table 3 also shows results for Topology C which significantly under-performed compared to Topology B. Since it exclusively consists of limb and joint TAFs without any spatial components, this makes keypoint localization and association rather difficult.

Topology B solves all of these problems elegantly. The longer cross-linked limb TAF connections preserve information even in the absence of motion or appearance of new people since the TAF effectively collapses to a PAF in such cases. This allows us to avoid association heuristics and makes the problem of new person identification trivial. With this representation, recurrence was observably beneficial due to true and consistent representation irrespective of magnitude of motion. As a side-advantage, this also allowed us to warm-start the TAF input with PAF providing more reliable initialization for tracking in the first frame.

For Model III, training beyond 5000 iterations gradually begins to harm the accuracy of the pose estimation resulting in reduced tracking performance as well. This is primarily due to the disparity in the amount of diverse data between COCO / MPII and PoseTrack datasets. For Model II, if we train on keypoints and PAFs modules and lock their weights afterwards, then follow with training only the TAF, this results in better performance with a significant boost in speed as well. Although Model I outperformed the other models with five stages for keypoints and PAFs; and a single recurrent stage for TAFs, however this comes at the expense of speed. Furthermore, we observe that an increase in mAP ends up sub-linearly increasing the MOTA as well.

#### Effect of Video Rate and Number of People on Tracking:

Finally, we performed a study on how the frame rate of the camera affects tracking accuracy, since a lower frame rate would require longer associations in pixel space.

We ran Lukas Kanade (LK) as a baseline tracker by replacing the TAF Module in Model I with LK ( $21 \times 21$  window size; 3 pyramid levels). Initially, we observe that there is roughly 2.0% improvement in MOTA as seen in

Method	Wrist-AP	Ankles-AP	mAP	MOTA	fps
Model I-A	56.2	56.4	66.0	58.5	14
Model I-B	56.3	56.9	66.3	59.4	13
Model II-A	54.9	53.0	64.4	57.4	28
Model II-B	55.0	53.5	64.6	58.4	27
Model III-B	51.9	49.5	61.6	57.8	30
Model III-C	42.5	40.5	55.2	49.9	36

Table 3: This table shows pose estimation and tracking performance for combinations of model types and topologies.




Figure 7: (a) This graph shows MOTA as a function of video frame rate for Temporal Affinity Fields (TAFs) and Lukas-Kanade (LK) tracker. The performance of TAFs is virtually invariant to frame rate or alternatively to the amount of motion between frames. (b) Our approach is effectively runtime-invariant to the number of people in the scene.

Fig. 7(a). However, we note that around 20% of the sequences have significant articulation and camera movement, where TAFs outperformed LK as the latter was not able to match keypoints across large displacements whereas TAFs found matches due to stronger descriptive power. TAFs were able to maintain tracking accuracy even with low frame-rate cameras, but with LK the MOTA drops off significantly (see Fig. 7(a)). Furthermore, Fig. 7(b) suggests that our approach is nearly runtime-invariant to number of people in the frame making it suitable for crowded scenes.

## 4.2. Comparison

We present results on PoseTrack dataset in Table 4 for 2017 validation set (top), 2017 test set (middle) and 2018 validation set (bottom). FlowTrack, JointFlow and PoseFlow are included as comparison in this table. FlowTrack is a top-down approach which means human detection is performed first followed by pose estimation. Due to this reason, it is significantly slower than bottom-up approaches such as ours. Model II-B with single scale is competitive with other bottom-up approaches while being 270% faster. However, multi-scale (MS) processing boosts performance by  $\sim 6\%$  and  $\sim 1.5\%$  for mAP and MOTA, respectively. We are also able to achieve competitive results on the PoseTrack 2018 Validation set while maintaining the best speeds amongst all reported results. Note that PoseTrack 2018 Test set was not released to public at the time of submission of this paper. Figure 8 shows some qualitative results.




Figure 8: Three example cases of tracking at  $\sim 30$  FPS on multiple targets. **Top / Middle:** Observe that tracking continues to function despite large motion displacements and occlusions. **Bottom:** A failure case where abrupt scene change causes ghosting, where previously tracked person appears in the new frame. This issue can be rectified through a warm-start.

## 5. Conclusion

In this paper, we first motivated recurrent Spatio-Temporal Affinity Fields (STAF) as the right approach for detection and tracking of articulated human pose in videos, especially for real-time reactive systems. We showed that leveraging the previous frame data within a recurrent structure and training on video sequences yields as good results as a multi-stage network albeit at much lower computation cost. We also demonstrated the stability of tracking accuracy at reduced frame rates for the TAF formulation, due to its ability to correlate keypoints over large pixel distances. This implies that our method can be deployed on low-power embedded systems which may not be able to run large networks at high frame rates, yet are able to maintain reasonable accuracy. Our new cross-linked limb temporal topology is able to generalize better than previous approaches due to strong associative power with PAF being a special case of TAF. We are also able to operate at the same consistent speed irrespective of the number of people due to bottom-up formulation. For future work, we plan to embed a re-identification module to handle cases of people leaving and reappearing in a camera view. Furthermore, detecting and triggering warm-start at every shot change has the potential to boost pose estimation and tracking performance.

**Acknowledgment:** Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DOI/IBC) contract number D17PC00340. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation/herein. Disclaimer: The views and conclusions contained herein

	Method	Wrist-AP	Ankles-AP	mAP	MOTA	fps
PoseTrack 2017 Validation						
Top-Down	Detect-and-track [9]	51.7	49.8	60.6	55.2	1.2
	FlowTrack - 152 [34]	72.4	67.1	76.7	65.4	-
	FlowTrack - 50 [34]	66.0	61.7	72.4	62.9	-
	MDPN - 152 [13]	77.5	71.4	80.7	66.0	-
Bottom-Up	PoseFlow [35]	61.1	61.3	66.5	58.3	10*
	JointFlow [8]	-	-	69.3	59.8	0.2
	Model II-B (SS)	55.0	53.5	64.6	58.4	27
	Model I-B (SS)	56.8	56.8	66.3	59.4	13
	<b>Model II-B (MS)</b>	<b>62.9</b>	<b>60.9</b>	<b>71.5</b>	<b>61.3</b>	<b>7</b>
	<b>Model I-B (MS)</b>	<b>65.0</b>	<b>62.7</b>	<b>72.6</b>	<b>62.7</b>	<b>2</b>
PoseTrack 2017 Testing						
Top-Down	Detect-and-track [9]	-	-	59.6	51.8	1.2
	Flowtrack - 152 [34]	70.7	64.9	73.9	57.6	-
	Flowtrack - 50 [34]	65.1	60.3	70.0	56.4	-
Bottom-Up	PoseTrack [2]	54.3	49.2	59.4	48.4	-
	BUTD [19]	52.9	42.6	59.1	50.6	-
	PoseFlow [35]	59.0	57.9	63.0	51.0	10*
	JointFlow [8]	53.1	50.4	63.3	53.1	0.2
	<b>Model II-B (MS)</b>	<b>62.8</b>	<b>59.5</b>	<b>69.6</b>	<b>52.4</b>	<b>7</b>
	<b>Model I-B (MS)</b>	<b>65.0</b>	<b>60.7</b>	<b>70.3</b>	<b>53.8</b>	<b>2</b>
PoseTrack 2018 Validation						
Bottom-Up	Model II-B (SS)	56.2	54.2	63.7	58.4	27
	Model I-B (SS)	58.3	56.7	64.9	59.6	13
	Model II-B (MS)	62.7	60.6	69.9	59.8	7
	Model I-B (MS)	64.7	62.0	70.4	60.9	3

Table 4: This table shows comparison on the PoseTrack datasets. For our approach, we report results with Models I / II and Top. B. The last column shows the speed in frames per second (\* excludes pose inference time). FlowTrack is a **top-down** approach using ResNet-152 (or 50); whereas JointFlow, PoseFlow and our approach are **bottom-up**.

are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.



## References

- [1] Posetrack leaderboard. <https://posetrack.net/leaderboard.php>. 2
- [2] M. Andriluka, U. Iqbal, E. Insafutdinov, L. Pishchulin, and A. Milan. PoseTrack: A benchmark for human pose estimation and tracking. In *CVPR*, 2018. 8
- [3] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2D human pose estimation: new benchmark and state of the art analysis. In *CVPR*, 2014. 1, 2
- [4] M. Andriluka, S. Roth, and B. Schiele. Monocular 3D pose estimation and tracking by detection. In *CVPR*, 2010. 2
- [5] Z. Cao, G. Hidalgo, T. Simon, S. Wei, and Y. Sheikh. Openpose: Realtime multi-person 2D pose estimation using part affinity fields. *CoRR*, abs/1812.08008, 2018. 5
- [6] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *CVPR*, 2017. 1, 2, 5
- [7] Y. Chen, Z. Wang, Y. Peng, Z. Zhang, G. Yu, and J. Sun. Cascaded pyramid network for multi-person pose estimation. *arXiv:1711.07319*, 2017. 2
- [8] A. Doering, U. Iqbal, and J. Gall. Joint flow: Temporal flow fields for multi person tracking. *CoRR*, abs/1805.04596, 2018. 2, 4, 7, 8
- [9] R. Girdhar, G. Gkioxari, L. Torresani, M. Paluri, and D. Tran. Detect-and-track: Efficient pose estimation in videos. *CoRR*, abs/1712.09184, 2017. 2, 8
- [10] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, and K. He. Detectron. <https://github.com/facebookresearch/detectron>, 2018. 2
- [11] G. Gkioxari, B. Hariharan, R. Girshick, and J. Malik. Using k-poselets for detecting people and localizing their keypoints. In *CVPR*, 2014. 2
- [12] R. A. Guler, N. Neverova, and I. Kokkinos. DensePose: Dense human pose estimation in the wild. In *CVPR*, 2018. 2
- [13] H. Guo, T. Tang, G. Luo, R. Chen, and Y. Lu. Multi-domain pose network for multi-person pose estimation and tracking. In *ECCV*, 2018. 8
- [14] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017. 2
- [15] E. Insafutdinov, M. Andriluka, L. Pishchulin, S. Tang, E. Levinkov, B. Andres, and B. Schiele. ArtTrack: Articulated multi-person tracking in the wild. In *CVPR*, 2017. 2
- [16] U. Iqbal and J. Gall. Multi-person pose estimation with local joint-to-person associations. In *ECCVW: Crowd Understanding*, 2016. 2
- [17] U. Iqbal, A. Milan, and J. Gall. PoseTrack: Joint multi-person pose estimation and tracking. In *CVPR*, 2017. 1, 2
- [18] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. J. Black. Towards understanding action recognition. In *ICCV*, 2013. 2
- [19] S. Jin, X. Ma, Z. Han, Y. Wu, and W. Yang. Towards multi-person pose tracking: Bottom-up and top-down methods. *ICCV*, 2017. 8
- [20] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, 2010. 2
- [21] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 1
- [22] 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. 2
- [23] Y. Luo, J. Ren, Z. Wang, W. Sun, J. Pan, J. Liu, J. Pang, and L. Lin. LSTM pose machines. In *CVPR*, 2018. 2
- [24] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *CoRR*, abs/1603.06937, 2016. 2
- [25] G. Papandreou, T. Zhu, L.-C. Chen, S. Gidaris, J. Tompson, and K. Murphy. PersonLab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. *arXiv:1803.08225*, 2018. 2
- [26] G. Papandreou, T. Zhu, N. Kanazawa, A. Toshev, J. Tompson, C. Bregler, and K. Murphy. Towards accurate multi-person pose estimation in the wild. In *CVPR*, 2017. 2
- [27] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. In *ICCV*, 2015. 2
- [28] L. Pishchulin, M. Andriluka, P. Gehler, and B. Schiele. Poselet conditioned pictorial structures. In *CVPR*, 2013. 2
- [29] L. Pishchulin, A. Jain, M. Andriluka, T. Thormahlen, and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR*, 2012. 2
- [30] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 3
- [31] J. Song, L. Wang, L. Van Gool, and O. Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *CVPR*, 2017. 2
- [32] M. Sun and S. Savarese. Articulated part-based model for joint object detection and pose estimation. In *ICCV*, 2011. 2
- [33] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016. 2
- [34] B. Xiao, H. Wu, and Y. Wei. Simple baselines for human pose estimation and tracking. *CoRR*, abs/1804.06208, 2018. 2, 8
- [35] Y. Xiu, J. Li, H. Wang, Y. Fang, and C. Lu. Pose flow: Efficient online pose tracking. In *BMVC*, 2018. 2, 8
- [36] Y. Yang and D. Ramanan. Articulated human detection with flexible mixtures of parts. *TPAMI*, 2013. 2
- [37] W. Zhang, M. Zhu, and K. Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *ICCV*, 2013. 2