

Explicit Spatiotemporal Joint Relation Learning for Tracking Human Pose

Xiao Sun
Microsoft Research
xias@microsoft.com

Chuankang Li *
Zhejiang University
lck-cad@zju.edu.cn

Stephen Lin
Microsoft Research
stevelin@microsoft.com

Abstract

We present a method for human pose tracking that is based on learning spatiotemporal relationships among joints. Beyond generating the heatmap of a joint in a given frame, our system also learns to predict the offset of the joint from a neighboring joint in the frame. Additionally, it is trained to predict the displacement of the joint from its position in the previous frame, in a manner that can account for possibly changing joint appearance, unlike optical flow. These relational cues in the spatial domain and temporal domain are inferred in a robust manner by attending only to relevant areas in the video frames. By explicitly learning and exploiting these joint relationships, our system achieves state-of-the-art performance on standard benchmarks for various pose tracking tasks including 3D body pose tracking in RGB video, 3D hand pose tracking in depth sequences, and 3D hand gesture tracking in RGB video.

1. Introduction

Human pose tracking is an essential component of many video understanding tasks such as visual surveillance, action recognition, and human-computer interaction. Compared to human pose estimation in single images, pose tracking in videos has shown to be a more challenging problem due in part to issues such as motion blur and uncommon poses with few or no examples in training sets. To help alleviate these problems, temporal information and spatial structure have been utilized in locating body joints.

A variety of approaches have been presented for exploiting temporal information in tracking human pose. Among them are techniques that use dense optical flow to propagate joint estimates from previous frames [31, 48, 52], local model fitting with the predicted pose of the preceding frame as initialization [11, 2, 38, 49, 45], and temporal smoothness priors as a pose estimation constraint [75, 41]. These methods take advantage of pose similarity from frame to frame, but do not learn explicitly about the effect of human mo-

tions on the appearance and displacement of joints. Rather, the motions are predicted through generic machinery like optical flow, local optimization of an initialized model, and smoothness constraints.

Other methods take advantage of the spatial structure of joints to facilitate joint localization in tracking. These include techniques that integrate graphical models into deep networks [60, 46, 69] and that employ structural priors [72, 11]. Such approaches are mainly designed to avoid invalid pose configurations rather than directly localize joints. Alternatively, several methods refine part locations based on the predicted positions of other parts [3, 66, 43, 23], which provides only vague constraints on joint positions.

In this paper, we propose an approach that *explicitly* learns about spatiotemporal joint relationships and utilizes them in conjunction with conventional joint heatmaps to elevate performance on human pose tracking. For learning temporal relationships, our system design reconfigures the optical flow framework to introduce elements from human pose estimation. Instead of densely predicting the motion vectors of general scene points as done for optical flow, our network is trained specifically to infer the displacements only of human joints, while also accounting for possible appearance changes due to motion. Inspired by heat maps in pose estimation, we leverage the other pixels to support the inference of each joint displacement. With training on labeled pose tracking datasets, our neural network learns to make use of information from both joint and non-joint pixels, as well as appearance change caused by joint motion, to predict joint displacements from one frame to the next.

Our system additionally learns spatial joint relationships, specifically the offset between a joint and a particular neighboring joint. These spatial relationships are learned in a manner analogous to the temporal relationships, except in the spatial domain. The spatial and temporal relationships, as well as single-frame joint heatmaps, all provide orthogonal cues based on different source information for predicting joint locations. We show that optimizing with respect to these cues together leads to appreciable improvements in pose tracking performance.

The presented technique comes with the following prac-

*Work done during internship at Microsoft Research Asia

tical benefits:

- *Generality*: This approach can be employed for various pose tracking applications (body/hand, 2D/3D, with/without an object) and with different input modalities (RGB/depth). In our experiments, the system is demonstrated on 3D human body pose tracking in RGB video, 3D hand gesture tracking from depth sequences, and 3D hand pose tracking in RGB video.
- *Flexibility*: The learned spatiotemporal relations are complementary not only to heatmap-based pose estimation, but also to other approaches such as local model fitting. Its speed-accuracy tradeoff is also controllable.
- *High performance*: State-of-the-art results are obtained on standard benchmarks for the aforementioned pose tracking applications, as shown in our experiments.

2. Related Work

Related research has exploited *joint structure* for single-frame human pose estimation and *temporal information* for pose tracking in video. We review these techniques, as well as methods for joint association within a multi-person context.

Joint Structure. Several methods incorporate structural constraints on joint locations by integrating graphical models [17] into deep networks. These graphical models have been used to remove anatomically incorrect outliers from part detections [60], constrain articulation with deformation models [46], and learn the spatial compatibility of neighboring parts to regularize the output [69]. The graphical constraints in these works are employed to enforce global pose consistency, rather than to provide explicit estimates of relative part positions.

In some other works, estimates of part locations are refined based on intermediate predictions of other parts. This refinement has been done by regressing part locations from the full set of part detection heatmaps [3], by iteratively computing the set of heatmaps within a multi-stage architecture [66, 43], and by sequentially predicting joint heatmaps using previously predicted parts in order of difficulty [23]. Refinement has also been performed in the feature maps of different joints via message passing [9]. In these methods, the intermediate part predictions provide only implicit constraints on joint localization.

Methods have also re-parameterized the pose representation to model joint structure. Such representations have been based on an overcomplete 3D pose dictionary [57] and kinematic models of the human body [73] and hand [74]. These re-parameterizations are either high-dimensional or

highly nonlinear, which makes optimization in deep networks hard. Alternatively, a geometric loss function has been introduced to enforce consistency of bone-length ratios in the human body [72], and this has been extended to include joint-angle limits and left-right symmetry [11]. Similar to the constraints of graphical models, these structural priors mainly serve to restrict pose predictions to valid configurations.

Recently, a practical pose re-parameterization based on bones instead of joints has been proposed together with a compositional loss [54]. We utilize these simple bone vectors to represent joint relationships in the spatial domain, and explicitly use bones as one of our learning targets. In [54], the performance is limited by its use of fully connected regression that maps the entire image feature to a local bone prediction. For better estimation, we instead propose to use a per-pixel regression together with an attention-like pixel weighting mechanism so as to take full advantage of the most relevant pixels while ignoring irrelevant ones.

Temporal Information. Many deep learning methods enhance pose estimation in the current frame with the help of dense optical flow. Among them are methods that concatenate the flow field with RGB features [31], compute flow-based alignment of heatmaps from previous frames to the current frame [48], and define temporal graph edges using flow [52]. Reliance on dense optical flow can be problematic, as flow estimates may be unreliable for non-rigid body regions that undergo appearance change during motion. In our work, we specifically train our model to predict the displacement only of human joints, with the support of information from relevant local areas. This task-specific approach provides greater reliability than generic flow, as shown empirically in Section 5.1. In addition, this approach is directly applicable to both 2D and 3D pose, while optical flow provides only 2D cues.

Temporal information is also exploited in techniques built on recurrent neural networks (RNNs) [33, 10, 25, 34], which account for information from previously predicted poses in determining the solution for the current frame. Similarly, some works fully connect the input and output pose sequences directly using a fully connected [11, 13] or convolutional [23] network, which draws global dependencies between the input and output. In contrast to these methods which implicitly encode temporal dependency into a network, we explicitly regress temporal joint displacements from only the most relevant input pixels. Attending only to information relevant to what is being processed has shown greater efficacy than recurrent computation in certain tasks [61].

Some works follow the *tracking by detection* paradigm, which performs detection in each frame, estimates the pose of each detection result with an off-the-shelf technique, then enforces spatio-temporal smoothness for tracking [75, 41].

Instead of employing generic smoothness priors, we seek a more targeted approach by specifically learning about human pose change.

Multi-Person Joint Association. Relationships between joints are also considered by methods for multi-person tracking and grouping. For the multi-person tracking problem, the goal is to estimate the pose of all persons appearing in the video and assign a unique identity to each person [29, 22, 67, 68, 15, 27, 12]. These approaches differ from one another in their choice of metrics and features used for measuring similarity between a pair of poses and in their algorithms for joint identity assignment. We note that in these techniques, the concept of ‘tracking’ is to associate joints of the same person together over time, using joints localized independently in each frame. By contrast, our work aims to improve joint localization by utilizing information from other frames.

Multi-person grouping aims to link the body parts of the same person in the context of multi-person pose estimation. To associate joints within an image, part affinity fields [4] and associative embedding tags [42] have been learned together with detection heatmaps. This information is then used to group body parts into separate identities. While these techniques learn joint relationships in terms of their association, our work instead focuses on explicitly predicting offsets between neighboring joints.

3. Explicit Spatiotemporal Joint Relation Learning

Given a video sequence, the pose tracking problem is to obtain the 2D (or 3D) position of all the K joints in all T frames, $\mathcal{J} = \{\mathbf{J}_k^t | k = 1, \dots, K; t = 1, \dots, T\}$. In solving this problem, we seek to supplement the per-joint heatmaps conventionally predicted from a single frame with **explicit** predictions of spatiotemporal joint relations. These relations include both the offset of a joint from a neighboring joint in the same frame and the offset of the joint from its position in the previous frame.

In modeling joint relations, we utilize a relation function $relation(k, t)$ that returns the index of the related joint for the k^{th} joint at time t , in either the spatial or temporal domain. This function is defined as

$$relation(k, t) = \begin{cases} (parent(k), t), & \text{in spatial domain} \\ (k, t-d), & \text{in temporal domain} \end{cases} \quad (1)$$

where $parent(k)$ returns the index of the parent joint for the k^{th} joint in a joint tree illustrated in Figure 1 (left).

The relationship between joint \mathbf{J}_k^t and its related joint $\mathbf{J}_{relation(k,t)}$ is defined as a directed vector pointing from the related joint to it:

$$\mathbf{R}_k^t = \mathbf{J}_k^t - \mathbf{J}_{relation(k,t)}. \quad (2)$$




Figure 1. Spatiotemporal joint relation learning problem.

	Joint	Related joint	Learning target: Relation	Input
General	\mathbf{J}_k^t	$\mathbf{J}_{relation(k,t)}$	$\mathbf{R}_k^t = \mathbf{J}_k^t - \mathbf{J}_{relation(k,t)}$	-
Spatial	\mathbf{J}_k^t	$\mathbf{J}_{parent(k)}^t$	$\mathbf{B}_k^t = \mathbf{J}_k^t - \mathbf{J}_{parent(k)}^t$	\mathbf{I}_t
Temporal	\mathbf{J}_k^t	\mathbf{J}_k^{t-d}	$\Delta_k^t = \mathbf{J}_k^t - \mathbf{J}_k^{t-d}$	$\mathbf{I}_{t-d}, \mathbf{I}_t$

Table 1. Notations.

Combining Eq. 2 and Eq. 1, we have

$$\mathbf{R}_k^t = \begin{cases} \mathbf{B}_k^t = \mathbf{J}_k^t - \mathbf{J}_{parent(k)}^t, & \text{for spatial relation} \\ \Delta_k^t = \mathbf{J}_k^t - \mathbf{J}_k^{t-d}, & \text{for temporal relation.} \end{cases} \quad (3)$$

The spatial joint relation \mathbf{B}_k^t is equivalent to the **bone vector** in Sun et al. [54], where the associated bone of the k^{th} joint is defined as the vector pointing from its parent to it. The temporal joint relation Δ_k^t represents the **joint displacement** from one time instance ($t-d$) to the next (t).

Given two successive frames \mathbf{I}_{t-d} and \mathbf{I}_t from a video sequence, our goal is to explicitly estimate the bone vector and joint displacement for each joint as shown in Figure 1 (right). Note that joint displacement involves the frames at both t and $t-d$. We thus stack \mathbf{I}_{t-d} and \mathbf{I}_t as input. On the other hand, the bone vector is computed from only a single input frame (\mathbf{I}_{t-d} or \mathbf{I}_t). The notations used in our spatiotemporal joint relation learning are summarized in Table 1.

3.1. Fully Connected Regression: A Baseline

A possible way to predict a relation \mathbf{R}_k^t is by using a CNN with a fully connected layer as done in [54] for predicting bone vectors. The input images are passed through convolutional layers to generate convolutional features. Following common practice [54, 5, 73, 72, 56], the features undergo spatial downsampling by average pooling. They are then given to a **fully connected** layer which outputs a $3K$ -dimensional bone or joint displacement vector. We refer to this holistic fully connected approach as the **Baseline**.

There is a clear drawback in the **Baseline** approach. Image features and outputs are fully connected, but most image areas/features are *irrelevant* to a specific relation. For

example, pixels that lie far from the joint’s vicinity are generally less informative for this prediction. The inclusion of non-relevant pixels and the over-emphasis of weakly relevant pixels would not only be unhelpful but could moreover be detrimental due to the noise they would introduce to the estimation. Worse still, the relevant image areas can change rapidly with respect to different joint locations, which would make the learning even more difficult.

3.2. Per-pixel Regression with Attention

To address the above issues, we present a method for per-pixel joint relation regression with an attention-like pixel weighting mechanism.

Joint Relation Map. Similar to heatmap based pose estimation, we generate a *separate* relation map M_k^t for each joint, where every pixel in a map predicts the relation R_k^t of the joint, as illustrated in Figure 2. For 3D pose, a joint relation map is a 3D vector field for the x, y, z dimensions. For compactness, the joint index k and time index t will be omitted in the rest of this section, as all joints at all times are processed separately and in the same way.

Pixel Weighting Map. However, not every pixel in a joint relation map is useful for estimating the relation of the joint. Therefore, an attention mechanism for filtering out these less informative pixel predictions is needed. To represent this information, a pixel weighting map W is generated for each joint relation map, as shown in Figure 2. Each pixel in W contains a weight value in $[0, 1]$ that indicates the relevance of the pixel for predicting the relation of this joint.

Intuitively, pixels closer to the related joint are generally more informative for estimating the joint relation. We denote *JointOne* as taking only one pixel at the ground truth location of the related joint $J_{relation(k,t)}$. For example, Figure 3 (middle row) shows the *JointOne* pixel sampling for (hip→) knee bone vector (yellow box) and hip joint displacement (red box) regression.

At the testing phase, the ground truth location of $J_{relation(k,t)}$ is not available. Instead, we use its prediction from the heat map estimation. To make the relation learning more robust to inaccurate heat map prediction, we propose a general pixel weighting formulation which is similar to pixel *translation augmentation* as shown in Figure 3 (bottom row), where box transparency indicates the weight decay of a pixel sample. Formally,

$$W = \mathcal{O}(F, \beta) \quad (4)$$

where \mathcal{O} is a decay function that determines the drop in weight for pixels farther away from $J_{relation(k,t)}$. F is a distance map in which each pixel value represents the distance from the pixel to $J_{relation(k,t)}$. β is the decay rate parameter. For the case of an exponential decay function,

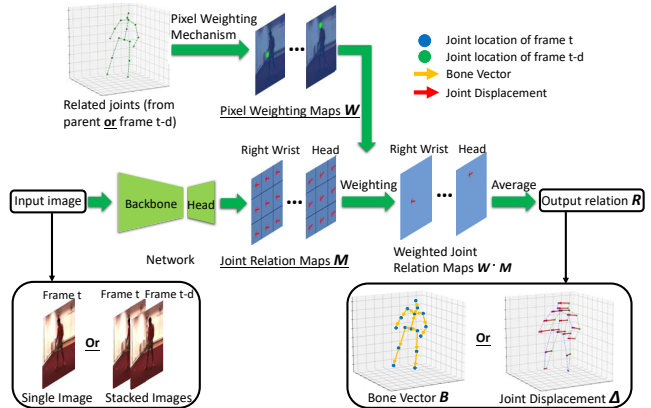


Figure 2. Overview of the explicit joint relation learning network.




Figure 3. Per-pixel regression on the most relevant pixels.

the pixel weighting map becomes

$$W = \mathcal{O}(F, \beta) = e^{-\beta F}. \quad (5)$$

Several decay functions are empirically investigated in Section 5.

Discussion. For a fully convolutional network, per-pixel regression essentially generates local image patches over each of the pixels using a sliding window and regresses the relation R from each local patch. The size of the sliding window is equal to the receptive field of the network. As shown in Figure 3, the per-pixel regression formulation combined with the pixel weighting mechanism is an effective way to filter out irrelevant pixel/patch samples and only preserve informative ones.

Figure 4 illustrates on a local patch the difference in information provided by heatmap based per-pixel joint classification and our explicit spatiotemporal joint relation regression. It can be seen that these three tasks approach the joint localization problem from orthogonal perspectives and using different observational data:

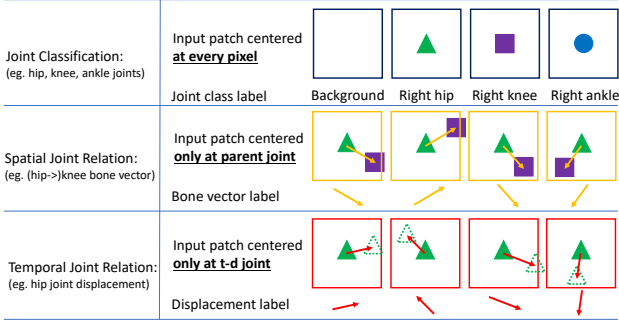


Figure 4. Orthogonal perspectives of joint classification, spatial joint relation regression, and temporal joint relation regression.

- **Joint Classification.** Determining whether a pixel is the correct type of joint from its local image appearance.
- **Spatial Joint Relation Regression.** Predicting a child joint position *conditioned* on its parent joint location.
- **Temporal Joint Relation Regression.** Estimating the change in joint location based on the *temporal transition*.

Our spatiotemporal relation learning effectively makes use of additional joint localization cues (dependence in space and continuity in time). The explicitly predicted bone vectors and joint displacements are thus *complementary* to single-frame per-joint detection results which are determined from only static joint appearance. Through integrating these cues, there exists a stronger foundation for inferring joint locations, via an optimization presented in Section 4.

Sparse Flow Baseline. Our explicit learning of temporal relationships is related to optical flow estimation, but it alters the conventional optical flow scheme to incorporate concepts from pose estimation. Rather than computing the motion of each pixel to form one dense flow map, our system focuses on predicting the motion only of pixels with specific semantics, namely joints. What’s more, a *separate* joint displacement map is generated for each joint, containing displacements predicted for that joint only. For an expanded comparison to optical flow, we introduce another baseline, called *SparseFlow*, that lies between optical flow and our method. It generates a *single* flow map as in optical flow but only predicts the motion of pixels at joint locations as in our method. A comprehensive comparison between optical flow, *SparseFlow* and our method is shown in Section 5.

3.3. Unified Network Architecture

We use a unified network architecture for both spatial and temporal joint relation learning as shown in Figure 2.

The only difference lies in the input. Spatial joint relation uses a single image as input while temporal joint relation stacks two consecutive images together as input. The input images are fed into a fully convolutional network. Then, a head network consisting of consecutive deconvolution layers is used to upsample the feature map to the targeted resolution as done in [56]. At the end, a 1x1 convolutional layer is used to produce the final joint relation maps.

Training Loss. We directly minimize the absolute difference (L1 norm) of the predicted and ground truth joint relation maps weighted by the pixel weighting map. This loss is expressed as

$$loss = \int_{\mathbf{p} \in \Omega} (\mathbf{W} \cdot \|\mathbf{M}^{pre} - \mathbf{M}^{gt}\|_1) \quad (6)$$

where Ω is the domain of the joint relation map. For other training details, please see the supplement.

Inference In the testing phase, the final prediction is the average of all predictions in a joint relation map weighted by the pixel weighting map:

$$\mathbf{R} = \frac{\int_{\mathbf{p} \in \Omega} (\mathbf{W} \cdot \mathbf{M}^{pre})}{\int_{\mathbf{p} \in \Omega} \mathbf{W}}. \quad (7)$$

4. Spatiotemporal Joint Relations for Tracking

Given single-frame per-joint pose predictions \mathbf{J}_k^t and spatiotemporal joint relation predictions \mathbf{B}_k^t and Δ_k^t , the pose tracking task can be formulated as a simple linear optimization problem. For a video sequence with T frames, we minimize the following error function to obtain the final tracking result $\hat{\mathbf{J}}_k^t$:

$$e = \sum_{t=1}^T \sum_{k=1}^K (\|\hat{\mathbf{J}}_k^t - \mathbf{J}_k^t\|^2 + \alpha * \|\hat{\mathbf{J}}_k^t - \hat{\mathbf{J}}_{parent(k)}^t - \mathbf{B}_k^t\|^2 + \sum_{n=1}^N (\gamma_n * \|\hat{\mathbf{J}}_k^t - \hat{\mathbf{J}}_k^{t-\mathbb{D}(n)} - \Delta_k^t\|^2)) \quad (8)$$

where α and γ are parameters to balance the single-frame and joint relation terms.

This optimization can be applied not only to two consecutive frames, but also to frames separated by different durations. Accounting for additional frames can potentially enhance the pose tracking. For this purpose, we define a set \mathbb{D} which represents N timesteps of duration d .

Minimizing e with respect to $\hat{\mathbf{J}}_k^t$ is a linear least squares problem with $(N+2)*K*T$ constraints, which is sufficient to solve the $K*T$ unknowns. The duration set \mathbb{D} can be arbitrary. Experimentally, we will consider four variants:

- $\mathbb{D}_f = \{1\}$, which is the common case of considering temporal joint relations only from the previous frame to the current frame (corresponding to one-step forward tracking).
- $\mathbb{D}_{fb} = \{1, -1\}$, which additionally considers displacements from the subsequent frame to the current frame (one-step forward and backward tracking).
- $\mathbb{D}_{mf} = \{1, 2, 3\}$, which accounts for temporal joint displacements from multiple past frames (multi-step forward tracking).
- $\mathbb{D}_{mfb} = \{1, 2, 3, -1, -2, -3\}$, which utilizes displacements from multiple past and future frames (multi-step forward and backward tracking).

5. Experiments


Our approach can be applied to various pose tracking scenarios, including human body or hand (w/ or w/o an object), RGB or depth input, and 2D or 3D output. We evaluate our approach extensively on four challenging datasets that cover all of these factors to show its generality. Specifically, 3D human pose dataset Human3.6M [28] of RGB images and 3D hand pose dataset MSRA Hand 2015 [55] of depth images are used in the main paper. 3D hand pose datasets Dexter+Object [53] and Stereo Hand Pose [71] of RGB-D images are used in the supplement.

5.1. Experiments on Human3.6M

Human3.6M [28] (HM36) is the largest 3D human pose benchmark of RGB images. It consists of 3.6 million video frames. For this benchmark, we employ the most widely used evaluation protocol in the literature [6, 59, 40, 76, 30, 37, 47, 70, 50, 2, 75, 58, 73, 56, 51]. Five subjects (S1, S5, S6, S7, S8, S9) are used in training and two subjects (S9, S11) are used in evaluation. The video sequences are typically downsampled for efficiency. We use three downsampling rates, namely no downsampling (25FPS), three-step downsampling (8FPS) and ten-step downsampling (2.5FPS). The lower the FPS, the larger the joint displacement between two consecutive frames.

For evaluation, most previous works use the mean per joint position error (*MPJPE*) in millimeters (mm). We call this metric *Joint Error*. To directly evaluate spatiotemporal joint relation performance, we use two additional metrics. The first metric is the mean per-bone position error, or *Bone Error* as in [54]. The second metric is the mean per-joint displacement error, or *Displ Error*, which adds the estimated joint displacement to the ground truth pose of the previous frame to obtain the current frame pose estimation and then evaluate its *Joint Error* in mm.

Ablation Study We first determine the best settings of our model for the explicit learning of spatiotemporal joint



Error (mm)	Full	RandHalf	RandOne	Baseline	Ours-JointOne
Bone Error	58.2	59.2	71.0	49.5	40.4 ↓18.4%
Displ Error	28.2	18.2	24.2	19.0	14.8 ↓22.1%

Table 2. Comparison of the fully connected *Baseline* and our per-pixel regression. Different pixel sampling strategies are investigated.

relations. For this purpose, we only use the joint relation evaluation metrics, namely *Bone Error* and *Displ Error*, because they isolate joint relation estimation from other factors such as the quality of single frame pose estimation and the specific tracking algorithm used.

Specifically, the fully connected regression *Baseline* and our attention-based per-pixel regression method are compared in Table 2. General optical flow (denoted as *DenseFlow*), the *SparseFlow* baseline, and our explicitly learned temporal joint relations (denoted as *Ours(Temp)*) are compared in Table 3. Different forms of the pixel weighting function for translation augmentation are investigated in Table 4 and Figure 5. We draw conclusions from these experiments one-by-one.

Per-pixel regression with relevant pixel sampling is superior to fully connected regression. This can be concluded from Table 2. Our per-pixel regression using only one pixel located at the prediction of the related joint (denoted as *Ours-JointOne*) outperforms the fully connected *Baseline* by a large margin, 18.4% and 22.1% relative improvement on spatial and temporal joint relations, respectively. Detailed *Bone Error* (left) and *Displ Error* (right) of *Baseline* and *Ours-JointOne* during training are shown in the top row.

For a more comprehensive understanding, we empirically evaluate several other pixel sampling strategies besides *Ours-JointOne*. *RandOne* randomly samples one pixel; *RandHalf* randomly samples half of the pixels, and *Full* uses all pixels. The results are reported in Table 2. A couple conclusions can be drawn. *First, sparse connections are superior to holistic mapping.* This can be seen from the lower performance of *Baseline* and *Full* compared to *Ours-JointOne*. *Second, pixel selection matters.* Randomly choosing the sparse connections (*RandOne*, *RandHalf*) yields lower performance than more purposeful selection (*Ours-JointOne*).

Explicitly learned temporal joint relation is superior to optical flow. Table 3 presents a comprehensive comparison

2.5FPS	Mean	Easy	Middle	Hard
DisplErr(mm)	(Max 39.2)	(62.5%)	(17.8%)	(19.7%)
DenseFlow	26.8	9.52	29.7	79.0
SparseFlow	22.2 _{↓17.2%}	7.89 _{↓17.1%}	27.3 _{↓8.08%}	62.7 _{↓20.6%}
Ours(Temp)	14.8 _{↓33.3%}	6.53 _{↓17.2%}	18.2 _{↓33.3%}	38.0 _{↓39.4%}
8FPS	Mean	Easy	Middle	Hard
DisplErr(mm)	(Max 12.7)	(88.9%)	(8.49%)	(2.65%)
DenseFlow	9.08	6.19	25.4	53.8
SparseFlow	7.85 _{↓13.5%}	5.36 _{↓13.4%}	22.7 _{↓10.6%}	44.0 _{↓18.2%}
Ours(Temp)	6.07 _{↓22.7%}	4.39 _{↓13.4%}	15.9 _{↓30.0%}	30.7 _{↓30.2%}
25FPS	Mean	Easy	Middle	Hard
DisplErr(mm)	(Max 4.32)	(99.1%)	(0.85%)	(0.02%)
DenseFlow	3.44	3.25	24.5	48.3
SparseFlow	3.05 _{↓11.3%}	2.87 _{↓11.7%}	22.3 _{↓8.98%}	46.1 _{↓4.55%}
Ours(Temp)	2.59 _{↓15.1%}	2.46 _{↓14.3%}	17.0 _{↓23.8%}	34.9 _{↓24.3%}

Table 3. Comparison to optical flow (*DenseFlow* and *SparseFlow*) for different frame rates. **Displ Err** is used. *Easy*, *Middle* and *Hard* represent joints whose displacement are less than 30mm, between 30mm and 60mm, and greater than 60mm, respectively.

between *DenseFlow*, *SparseFlow* and *Ours(Temp)*. *SparseFlow* and *Ours(Temp)* are trained using the same network setting and data for a fair comparison. For *DenseFlow*, we use the state-of-the-art *FlowNet2.0* [26] to compute optical flow between two consecutive frames, and then use the same pixel weighting scheme as in *Ours(Temp)*. Since joint displacement in the depth dimension cannot be determined from *DenseFlow*, we directly use the corresponding displacement in depth from *Ours(Temp)* for *DenseFlow*. The *Displ Error* metric is used for evaluation.

Three conclusions can be drawn from Table 3. *First*, the pixel weighting mechanism that only uses relevant pixels is effective. This can be concluded from the *SparseFlow* baseline outperforming *DenseFlow* in all the cases, where the relative performance gain is shown as subscripts. *Second*, the separate joint relation map for each joint to represent a more explicit joint semantic is effective. This can be concluded from *Ours(Temp)* further improving *SparseFlow* in all the cases by a large margin, where the relative performance gain to *SparseFlow* is shown as subscripts. *Third*, *Ours(Temp)* is especially superior under large joint displacement. This can be seen from the larger relative improvement under lower frame rates (*Ours(Temp)* outperforms *SparseFlow* by 33.3% at 2.5FPS, but by 15.1% at 25FPS) and for more challenging joints with large displacement (*Ours(Temp)* outperforms *SparseFlow* by 39.4% in *Hard* level, but by 17.2% in *Easy* level at 2.5FPS).

Decay function for translation augmentation is effective. We investigate four forms of decay function $\mathcal{O}(\mathbf{J}, \beta)$ in Eq. 4. They are *Binary*, *Gaussian*, *Linear* and *Exponential* defined in Table 4. When $\beta = 0$ or ∞ , $\mathcal{O}(\mathbf{J}, \beta)$ degenerates to *Ours-JointOne* and *Full*, respectively. It is seen that when β is slightly larger than 0, we obtain better performance ($\beta = 5$ in the *Binary* case). Similar observations

function	formulation
Binary	$\begin{cases} 1 & \text{if } \mathbf{J} \leq \beta \\ 0 & \text{if } \mathbf{J} > \beta \end{cases}$
Gaussian	$e^{-\beta \mathbf{J}^2}$
Linear	$1 - \beta \mathbf{J}$
Exponential	$e^{-\beta \mathbf{J}}$

Table 4. Different definitions of the decay function $\mathcal{O}(\mathbf{J}, \beta)$ for Eq. 4.




Figure 5. Performance for the *Binary* decay function with respect to β .

Joint Error (mm)	HM36	HM36MPII	description
Integral Reg [56]	63.8	49.5	single frame baseline
Ours-F	61.7	48.8	+ \mathbb{D}_f
Ours-FB	61.1	48.5	+ \mathbb{D}_{fb}
Ours-MF	59.6	47.8	+ \mathbb{D}_{mf}
Ours-Temporal	59.3 _{↓7.1%}	47.7 _{↓3.6%}	+ \mathbb{D}_{mfb}
Ours-Spatial	62.1 _{↓2.7%}	47.4 _{↓4.2%}	+bone vector
Ours-Spatiotemporal	58.1 _{↓8.9%}	45.5 _{↓8.1%}	+ \mathbb{D}_{mfb} +bone vector

Table 5. Effects of spatiotemporal joint relations on tracking. (See supplement for detailed results on all joints and dimensions.)

are obtained using other decay functions. The performance differences between different decay functions are minor, but an ensemble of models with different functions yields considerable improvement. We thus use this *Ours-Ensemble* model as our final joint relation estimator for tracking.

Effects of Spatiotemporal Joint Relations on Tracking. Now we apply the explicitly learned spatiotemporal joint relations to the tracking framework introduced in Section 4 for better joint localization in videos. Hence, the joint localization metric *Joint Error* is used for evaluation. For the single frame pose estimator, we use the Pytorch implementation [1] of Integral Regression [56]. Specifically, two single-frame baselines are used. The first uses only *HM36* data for training and employs a two-stage network architecture. The second baseline mixes *HM36* and *MPII* data for training and uses a one-stage network architecture (*HM36MPII*). Note that these two baselines are the state of the art and thus produce strong baseline results.

Table 5 shows how joint localization accuracy improves using our explicitly learned spatiotemporal joint relations. The *Joint Error* metric and 8FPS frame rates are used for evaluation. We can draw three conclusions. First, multi-step tracking is effective. With more elements in the duration set, better performance is obtained ($\mathbb{D}_{mfb} > \mathbb{D}_{mf} > \mathbb{D}_{fb} > \mathbb{D}_f$). In other words, the more observations we get from different durations in Eq. 8, the better pose tracking result we can get. Second, both *Ours-Spatial* (or *Ours-S*) and *Ours-Temporal* (or *Ours-T*) effectively improve all single-frame baselines. Larger relative improvement can

Method	Tome [59]	Moreno [40]	Zhou [76]	Jahangiri [30]	Mehta [37]	Martinez [36]	Kanazawa [32]	Fang [16]	Sun [54]	Sárándi [51] [†]	Sun [56]	Dabral [11]*	Hossain [25]*	Ours *
MPJPE	88.4	87.3	79.9	77.6	72.9	62.9	88.0	60.4	59.1	54.2	49.6	<u>52.1</u>	<u>51.9</u>	45.5

Table 6. Comparisons to previous work on Human3.6M. All methods used extra 2D training data. Methods with * exploit temporal information. Ours gives the lowest error and improves on the state of the art [56] by 4.1mm (relative 8.3%).

Method	Zhou [75]	Tekin [58]	Xingyi [73]	Sun [54]	Pavlakos [47]	Sun [56]	Lin [33]*	Coskun [10]*	Ours*
MPJPE	113.0	125.0	107.3	92.4	71.9	64.1	73.1	71.0	58.1

Table 7. Comparison to previous work on Human3.6M. No extra training data is used. Methods with * exploit temporal information. Ours yields the lowest error and improves on the state of the art [56] by 6.0mm (relative 9.4%).

be obtained on the low-performance single-frame baseline. Third, combining both spatial and temporal relations (*Ours-Spatiotemporal* or *Ours-ST*) leads to further improvements and performs the best.

Comparison with the state of the art Previous works are commonly divided into two categories. The first uses extra 2D data for training. The second only uses *HM36* data for training. They are compared to our method in Table 6 and Table 7, respectively. Methods marked with * are tracking based methods that exploit temporal information. Note that in Table 6, although Sárándi et al. [51][†] do not use extra 2D pose ground truth, they use extra 2D images for synthetic occlusion data augmentation where the occluders are realistic object segments from the Pascal VOC 2012 dataset [14]. Their data augmentation technique is also complementary to our method.

Our method sets a new state of the art on the Human3.6M benchmark under both categories. Specifically, it improves the state-of-the-art by 4.1mm (relative 8.3%) in Table 6, and 6.0mm (relative 9.4%) in Table 7.

5.2. Experiments on MSRA15

MSRA Hand 2015 (MSRA15) [55] is a standard benchmark for 3D hand pose estimation from depth. It consists of 76.5k video frames in total. For evaluation, we use two common accuracy metrics in the literature [55, 64, 7, 20, 24, 44, 62, 63]. The first one is mean 3D *Joint Error*. The second one is the percentage of correct frames (*PCF*). A frame is considered correct if its maximum joint error is less than a small threshold [55]. For single-frame pose estimation, many recent works provide their single-frame predictions [63, 8, 21, 18]. We use all of them as our single-frame baselines. In addition, we re-implement a new single-frame baseline using Integral Regression [56].

Effect of Spatiotemporal Joint Relations on Tracking

We apply our explicitly learned joint relations on all of the single-frame baselines for tracking, and list the results in Table 8. The mean 3D joint error is used as the evaluation metric. Our observations are similar to or even stronger than those for the *HM36* experiments.

First, our spatial and/or temporal joint relations signifi-

Joint Error (mm)	IntReg [56]
single frame	8.42
Ours-F	7.10
Ours-FB	6.78
Ours-MF	6.42
Ours-T	6.35 _{↓24.6%}
Ours-S	6.67 _{↓20.8%}
Ours-ST	5.96 _{↓29.2%}

Joint Err (mm)	PointNet [18]	SHPR-Net [8]	Point2Point [21]	DenseReg [63]
single	8.50	7.76	7.71	7.23
Ours-T	7.19 _{↓15.4%}	7.02 _{↓9.54%}	6.80 _{↓11.8%}	6.96 _{↓3.73%}
Ours-S	7.54 _{↓11.3%}	7.09 _{↓8.63%}	6.89 _{↓10.6%}	6.45 _{↓10.8%}
Ours-ST	6.68 _{↓21.4%}	6.56 _{↓15.5%}	6.28 _{↓18.5%}	6.44 _{↓10.9%}

Table 8. Effects of spatiotemporal joint relations on tracking. The graph shows a comparison to IntReg [56] on the *PCF* metric.

Methods	Sun [55]	Ge [19]	Madadi [35]	Wan [62]
Joint Error	15.2mm	13.2mm	12.8mm	12.2mm
Wang [65]	Ge [20]	Oberweger [44]	Chen [7]	Ge [18]
9.79mm	9.58mm	9.50mm	8.65mm	8.50mm
Chen [8]	Ge [21]	Moon [39]	Wan [63]	Ours
7.76mm	7.71mm	7.49mm	7.23mm	5.96mm

Table 9. Comparison to previous work on MSRA15. Ours yields the lowest error and improves [63] by 1.27mm (relative 17.6%).

cantly improve all single-frame baselines. The relative performance gains over the single-frame baselines are shown in the subscripts. Second, the improvement is especially large on the per-pixel joint classification baseline [56]. Specifically, our method improves the Integral Regression baseline by 2.46mm (relative 29.2%), establishing a new record of 5.96mm mean 3D joint error on this benchmark.

Comparisons to the state of the art We compare our method to previous works in Table 9 with the mean 3D joint error metric. Our method achieves a 5.96mm mean 3D joint error and outperforms the previous state-of-the-art result of 7.23mm [63] by a large margin (17.6% relative improvement). Please see the supplement for more comparisons using the *PCF* metric.

References

- [1] Integral Regression Pytorch Github. <https://github.com/JimmySuen/integral-human-pose>. 7
- [2] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *European Conference on Computer Vision*, pages 561–578. Springer, 2016. 1, 6
- [3] A. Bulat and G. Tzimiropoulos. Human pose estimation via convolutional part heatmap regression. In *European Conference on Computer Vision*, pages 717–732. Springer, 2016. 1, 2
- [4] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. *arXiv preprint arXiv:1611.08050*, 2016. 3
- [5] J. Carreira, P. Agrawal, K. Fragkiadaki, and J. Malik. Human pose estimation with iterative error feedback. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4733–4742, 2016. 3
- [6] C.-H. Chen and D. Ramanan. 3d human pose estimation= 2d pose estimation+ matching. *arXiv preprint arXiv:1612.06524*, 2016. 6
- [7] X. Chen, G. Wang, H. Guo, and C. Zhang. Pose guided structured region ensemble network for cascaded hand pose estimation. *arXiv preprint arXiv:1708.03416*, 2017. 8
- [8] X. Chen, G. Wang, C. Zhang, T.-K. Kim, and X. Ji. Shprnet: Deep semantic hand pose regression from point clouds. *IEEE Access*, 6:43425–43439, 2018. 8
- [9] X. Chu, W. Ouyang, H. Li, and X. Wang. Structured feature learning for pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4715–4723, 2016. 2
- [10] H. Coskun, F. Achilles, R. S. DiPietro, N. Navab, and F. Tombari. Long short-term memory kalman filters: Recurrent neural estimators for pose regularization. In *ICCV*, pages 5525–5533, 2017. 2, 8
- [11] R. Dabral, A. Mundhada, U. Kusupati, S. Afaque, A. Sharma, and A. Jain. Learning 3d human pose from structure and motion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 668–683, 2018. 1, 2, 8
- [12] A. Doering, U. Iqbal, and J. Gall. Joint flow: Temporal flow fields for multi person tracking. *arXiv preprint arXiv:1805.04596*, 2018. 3
- [13] M. Einfalt, D. Zecha, and R. Lienhart. Activity-conditioned continuous human pose estimation for performance analysis of athletes using the example of swimming. *arXiv preprint arXiv:1802.00634*, 2018. 2
- [14] M. Everingham and J. Winn. The pascal visual object classes challenge 2012 (voc2012) development kit. *Pattern Analysis, Statistical Modelling and Computational Learning, Tech. Rep*, 2011. 8
- [15] M. Fabbri, F. Lanzi, S. Calderara, A. Palazzi, R. Vezzani, and R. Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. *arXiv preprint arXiv:1803.08319*, 2018. 3
- [16] H.-S. Fang, Y. Xu, W. Wang, X. Liu, and S.-C. Zhu. Learning pose grammar to encode human body configuration for 3d pose estimation. In *AAAI*, 2018. 8
- [17] P. F. Felzenszwalb and D. P. Huttenlocher. Pictorial structures for object recognition. *International Journal of Computer Vision*, 61(1):55–79, 2005. 2
- [18] L. Ge, Y. Cai, J. Weng, and J. Yuan. Hand pointnet: 3d hand pose estimation using point sets. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8417–8426, 2018. 8
- [19] L. Ge, H. Liang, J. Yuan, and D. Thalmann. Robust 3d hand pose estimation in single depth images: from single-view cnn to multi-view cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3593–3601, 2016. 8
- [20] L. Ge, H. Liang, J. Yuan, and D. Thalmann. 3d convolutional neural networks for efficient and robust hand pose estimation from single depth images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, page 5, 2017. 8
- [21] L. Ge, Z. Ren, and J. Yuan. Point-to-point regression pointnet for 3d hand pose estimation. *ECCV, Springer*, 1, 2018. 8
- [22] R. Girdhar, G. Gkioxari, L. Torresani, M. Paluri, and D. Tran. Detect-and-track: Efficient pose estimation in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 350–359, 2018. 3
- [23] G. Gkioxari, A. Toshev, and N. Jaitly. Chained predictions using convolutional neural networks. In *European Conference on Computer Vision*, pages 728–743. Springer, 2016. 1, 2
- [24] H. Guo, G. Wang, X. Chen, C. Zhang, F. Qiao, and H. Yang. Region ensemble network: Improving convolutional network for hand pose estimation. In *Image Processing (ICIP), 2017 IEEE International Conference on*, pages 4512–4516. IEEE, 2017. 8
- [25] M. R. I. Hossain and J. J. Little. Exploiting temporal information for 3d pose estimation. *arXiv preprint arXiv:1711.08585*, 2017. 2, 8
- [26] 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 *IEEE conference on computer vision and pattern recognition (CVPR)*, volume 2, page 6, 2017. 7
- [27] E. Insafutdinov, M. Andriluka, L. Pishchulin, S. Tang, E. Levinkov, B. Andres, and B. Schiele. Arttrack: Articulated multi-person tracking in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 4327. IEEE, 2017. 3
- [28] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2014. 6
- [29] U. Iqbal, A. Milan, and J. Gall. Posetrack: Joint multi-person pose estimation and tracking. *arXiv preprint arXiv:1611.07727*, 2016. 3

- [30] E. Jahangiri and A. L. Yuille. Generating multiple hypotheses for human 3d pose consistent with 2d joint detections. *arXiv preprint arXiv:1702.02258*, 2017. 6, 8
- [31] A. Jain, J. Tompson, Y. LeCun, and C. Bregler. Modeep: A deep learning framework using motion features for human pose estimation. In *Asian conference on computer vision*, pages 302–315. Springer, 2014. 1, 2
- [32] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. *arXiv preprint arXiv:1712.06584*, 2017. 8
- [33] M. Lin, L. Lin, X. Liang, K. Wang, and H. Cheng. Recurrent 3d pose sequence machines. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5543–5552. IEEE, 2017. 2, 8
- [34] Y. Luo, J. Ren, Z. Wang, W. Sun, J. Pan, J. Liu, J. Pang, and L. Lin. Lstm pose machines. *arXiv preprint arXiv:1712.06316*, 2017. 2
- [35] M. Madadi, S. Escalera, A. Carruesco, C. Andujar, X. Baró, and J. González. Occlusion aware hand pose recovery from sequences of depth images. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pages 230–237. IEEE, 2017. 8
- [36] J. Martinez, R. Hossain, J. Romero, and J. J. Little. A simple yet effective baseline for 3d human pose estimation. *arXiv preprint arXiv:1705.03098*, 2017. 8
- [37] D. Mehta, H. Rhodin, D. Casas, O. Sotnychenko, W. Xu, and C. Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. *arXiv preprint arXiv:1611.09813*, 2016. 6, 8
- [38] D. Mehta, S. Sridhar, O. Sotnychenko, H. Rhodin, M. Shafiei, H.-P. Seidel, W. Xu, D. Casas, and C. Theobalt. Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 36(4):44, 2017. 1
- [39] G. Moon, J. Y. Chang, and K. M. Lee. V2v-posesnet: Voxel-to-voxel prediction network for accurate 3d hand and human pose estimation from a single depth map. In *CVPR*, volume 2, 2018. 8
- [40] F. Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. *arXiv preprint arXiv:1611.09010*, 2016. 6, 8
- [41] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. Gnerated hands for real-time 3d hand tracking from monocular rgb. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 49–59, 2018. 1, 2
- [42] A. Newell, Z. Huang, and J. Deng. Associative embedding: End-to-end learning for joint detection and grouping. In *Advances in Neural Information Processing Systems*, pages 2277–2287, 2017. 3
- [43] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016. 1, 2
- [44] M. Oberweger and V. Lepetit. Deepprior++: Improving fast and accurate 3d hand pose estimation. In *ICCV workshop*, volume 840, page 2, 2017. 8
- [45] I. Oikonomidis, N. Kyriazis, and A. A. Argyros. Efficient model-based 3d tracking of hand articulations using kinect. In *BmVC*, volume 1, page 3, 2011. 1
- [46] W. Ouyang and X. Wang. Joint deep learning for pedestrian detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2056–2063, 2013. 1, 2
- [47] G. Pavlakos, X. Zhou, K. G. Derpanis, and K. Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. *arXiv preprint arXiv:1611.07828*, 2016. 6, 8
- [48] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1913–1921, 2015. 1, 2
- [49] C. Qian, X. Sun, Y. Wei, X. Tang, and J. Sun. Realtime and robust hand tracking from depth. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1106–1113, 2014. 1
- [50] G. Rogez and C. Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *Advances in Neural Information Processing Systems*, pages 3108–3116, 2016. 6
- [51] I. Sáráandi, T. Linder, K. O. Arras, and B. Leibe. Synthetic occlusion augmentation with volumetric heatmaps for the 2018 eccv posetrack challenge on 3d human pose estimation. *arXiv preprint arXiv:1809.04987*, 2018. 6, 8
- [52] J. Song, L. Wang, L. Van Gool, and O. Hilliges. Thin-slicing network: A deep structured model for pose estimation in videos. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2017. 1, 2
- [53] S. Sridhar, F. Mueller, M. Zollhöfer, D. Casas, A. Oulasvirta, and C. Theobalt. Real-time joint tracking of a hand manipulating an object from rgb-d input. In *European Conference on Computer Vision*, pages 294–310. Springer, 2016. 6
- [54] X. Sun, J. Shang, S. Liang, and Y. Wei. Compositional human pose regression. In *International Conference on Computer Vision*, 2017. 2, 3, 6, 8
- [55] X. Sun, Y. Wei, S. Liang, X. Tang, and J. Sun. Cascaded hand pose regression. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 824–832, 2015. 6, 8
- [56] X. Sun, B. Xiao, F. Wei, S. Liang, and Y. Wei. Integral human pose regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 529–545, 2018. 3, 5, 6, 7, 8
- [57] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. *arXiv preprint arXiv:1605.05180*, 2016. 2
- [58] B. Tekin, A. Rozantsev, V. Lepetit, and P. Fua. Direct prediction of 3d body poses from motion compensated sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 991–1000, 2016. 6, 8
- [59] D. Tome, C. Russell, and L. Agapito. Lifting from the deep: Convolutional 3d pose estimation from a single image. *arXiv preprint arXiv:1701.00295*, 2017. 6, 8
- [60] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *Advances in neural information processing systems*, pages 1799–1807, 2014. 1, 2

- [61] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017. 2
- [62] C. Wan, T. Probst, L. Van Gool, and A. Yao. Crossing nets: Combining gans and vaes with a shared latent space for hand pose estimation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017. 8
- [63] C. Wan, T. Probst, L. Van Gool, and A. Yao. Dense 3d regression for hand pose estimation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.(CVPR)*, pages 1–10, 2018. 8
- [64] C. Wan, A. Yao, and L. Van Gool. Hand pose estimation from local surface normals. In *European conference on computer vision*, pages 554–569. Springer, 2016. 8
- [65] G. Wang, X. Chen, H. Guo, and C. Zhang. Region ensemble network: Towards good practices for deep 3d hand pose estimation. *Journal of Visual Communication and Image Representation*, 2018. 8
- [66] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016. 1, 2
- [67] B. Xiao, H. Wu, and Y. Wei. Simple baselines for human pose estimation and tracking. *arXiv preprint arXiv:1804.06208*, 2018. 3
- [68] Y. Xiu, J. Li, H. Wang, Y. Fang, and C. Lu. Pose flow: Efficient online pose tracking. *arXiv preprint arXiv:1802.00977*, 2018. 3
- [69] W. Yang, W. Ouyang, H. Li, and X. Wang. End-to-end learning of deformable mixture of parts and deep convolutional neural networks for human pose estimation. In *CVPR*, 2016. 1, 2
- [70] H. Yasin, U. Iqbal, B. Kruger, A. Weber, and J. Gall. A dual-source approach for 3d pose estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4948–4956, 2016. 6
- [71] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang. 3d hand pose tracking and estimation using stereo matching. *arXiv preprint arXiv:1610.07214*, 2016. 6
- [72] X. Zhou, Q. Huang, X. Sun, X. Xue, and Y. Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *International Conference on Computer Vision*, 2017. 1, 2, 3
- [73] X. Zhou, X. Sun, W. Zhang, S. Liang, and Y. Wei. Deep kinematic pose regression. In *Computer Vision—ECCV 2016 Workshops*, pages 186–201. Springer, 2016. 2, 3, 6, 8
- [74] X. Zhou, Q. Wan, W. Zhang, X. Xue, and Y. Wei. Model-based deep hand pose estimation. *arXiv preprint arXiv:1606.06854*, 2016. 2
- [75] X. Zhou, M. Zhu, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4966–4975, 2016. 1, 2, 6, 8
- [76] X. Zhou, M. Zhu, G. Pavlakos, S. Leonardos, K. G. Derpanis, and K. Daniilidis. Monocap: Monocular human motion capture using a cnn coupled with a geometric prior. *arXiv preprint arXiv:1701.02354*, 2017. 6, 8