Predicting People's 3D Poses from Short Sequences

Bugra Tekin[a] Xiaolu Sun[a] Xinchao Wang[a] Vincent Lepetit[a,b] Pascal Fua[a]
[a]Computer Vision Laboratory, École Polytechnique Fédérale de Lausanne (EPFL)
[b]Institute for Computer Graphics and Vision, Graz University of Technology
{bugra.tekin, xiaolu.sun, xinchao.wang, pascal.fua}@epfl.ch, lepetit@icg.tugraz.at

Abstract

We propose an efficient approach to exploiting motion information from consecutive frames of a video sequence to recover the 3D pose of people. Instead of computing candidate poses in individual frames and then linking them, as is often done, we regress directly from a spatio-temporal block of frames to a 3D pose in the central one. We will demonstrate that this approach allows us to effectively overcome ambiguities and to improve upon the state-of-the-art on challenging sequences.

1. Introduction

In recent years, impressive motion capture results have been demonstrated using depth cameras but 3D body pose recovery from ordinary video sequences remains extremely challenging. Nevertheless, there is great interest in doing so, both because cameras are becoming ever cheaper and more prevalent and because there are so many potential applications. These include athletic training, surveillance, entertainment, and electronic publishing.

Most early approaches to monocular 3D pose tracking involved recursive frame-to-frame tracking and were found to be brittle, due to distractions and occlusions from other people or objects in the scene. Since then, the focus has shifted to "tracking by detection," which involves detecting human pose more or less independently in every frame followed by linking detection across frames [26, 2], which is much more robust to algorithmic failures in isolated frames. In such approaches, motion information is exploited only *a posteriori* by the linking procedure. It essentially eliminates erroneous poses by selecting compatible candidates over consecutive frames. If there is none or few correct ones among them, nothing can be done.

Recently, [16] proposed an effective single-frame approach by learning a regressor from a kernel embedding of 2D HoG features to 3D poses. Since then, excellent results have also been reported using a Convolutional Neural Net [22]. Even for such state-of-the-art approaches depending on 2D appearance information, there still exist many detection errors caused by inherent ambiguities of the projection from 3D to 2D, including self-occlusion and mirroring as depicted in Fig. 1. When such errors happen frequently for several frames in a row, enforcing temporal consistency after the fact is not enough.

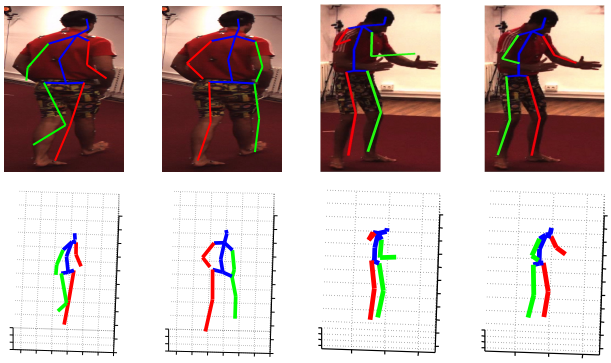

(a) $e^{\chi 2}$-HoG (b) Our (c) $e^{\chi 2}$-HoG (d) Our

Figure 1. 3D pose estimation in the Human3.6m dataset. The recovered 3D skeletons are reprojected into the images in the top row and shown by themselves in the bottom one. (a,c) Results obtained using a single-frame [16] are penalized by self-occlusions and mirror ambiguities. (b,d) By contrast, our approach can reliably recover 3D poses in such cases by collecting appearance and motion evidence from multiple frames simultaneously. Best viewed in color.

In this paper, we show that we can overcome these limitations by using appearance and motion information simultaneously and regressing directly from short sequences of frames to 3D poses in the central one. To this end, we find bounding boxes around the potential subjects, shift them so that the person inside them remains centered, compute 3D HoG features, and learn a mapping from these features to 3D poses. This prevents errors from which the above-mentioned methods cannot recover and we will show that it yields substantial overall performance increases over state-of-the-art methods on challenging test datasets.

Spatiotemporal image cubes have been used before for

action recognition purposes [20, 39], person detection [24], and 2D pose estimation [11] but, to the best of our knowledge, not for discriminative 3D human pose estimation. The only such method we know of is that of [40] which is computationally very expensive. Besides, it is not discriminative, and its 3D accuracy has not been reported.

The contribution of this paper is therefore a novel approach to combining appearance and motion clues very early in the 3D pose estimation process. Because we regress directly from a spatio-temporal volume to a 3D pose, we have little need for complex uncertainty estimation and propagation schemes to handle the ambiguities of the human pose and achieve superior performance. However, as we will show, the ability to properly align successive image windows from consecutive frames is key to obtaining it.

In the remainder of the paper, we first review related work in Section 2. Then, in Section 3, we describe our method. Finally, in Section 4, we present the results obtained on challenging datasets and show that we improve upon the state-of-the-art.

2. Related Work

Approaches to estimating the 3D human pose can be classified into two main categories, depending on whether they rely on still images or image sequences. We briefly review both kinds below. In the results section, we will demonstrate that we outperform state-of-the-art representatives of each.

3D Human Pose Estimation in Single Images. Early approaches tended to rely on generative models to search the state space for a plausible configuration of the skeleton that would align with the image evidence [9, 35, 32, 38, 12]. These methods remain competitive provided that a good enough initialization can be supplied. More recent ones [6, 3] extend 2D pictorial structure approaches [10] into the 3D domain. However, in addition to their high computational cost, they tend to have difficulty localizing people's arms accurately because the corresponding appearance cues are weak and easily confused with the background [27].

By contrast, discriminative regression-based approaches [1, 34, 30, 4, 15] build a direct mapping from image evidence to 3D poses. They have been shown to be effective, especially if a large training dataset [16] is available. Within this context, rich features encoding depth [29] and body part information [15, 22] have been shown to be effective at increasing the estimation accuracy.

3D Human Pose Estimation in Image Sequences. Such approaches also fall into two main classes.

The first class involves frame-to-frame tracking and dynamical models that rely on Markov dependencies between consecutive frames. Their main weakness is that they require initialization and cannot recover from tracking failures.

To address these shortcomings, the second class focuses on detecting candidate poses in individual frames followed by linking them across frames in a temporally consistent manner. For example, in [2], initial pose estimates are refined using 2D tracklet-based estimates. In [28], the pairwise relationships of joints within and between frames are modeled by an ensemble of tractable submodels. [13] integrates single-frame pose recovery with K-best trajectories and model texture adaptation. In [41], dense optical flow is used to link articulated shape models in adjacent frames. Non-maxima suppression is then employed to merge pose estimates across frames in [7]. By contrast to these approaches, we capture the temporal information earlier in the process by extracting spatiotemporal features from image cubes of short sequences and regressing to 3D poses.

While they have long been used for action recognition [20, 39], person detection [24], and 2D pose estimation [11], spatiotemporal features have been underused for 3D body pose estimation purposes. The only recent approach we are aware of is that of [40] that involves building a set of point trajectories corresponding to high joint responses and matching them to motion capture data. One drawback of this approach is its very high computational cost. Also, while the 2D results look promising, no quantitative 3D results are provided in the paper and no code is available for comparison purposes.

3. Method

Our approach involves finding bounding boxes around people in consecutive frames, aligning these bounding boxes to form spatiotemporal volumes, and learning a mapping from these volumes to a 3D pose in their central frame.

In the remainder of this section, we first introduce our formalism and then describe each individual step, as depicted by Fig. 2.

3.1. Formalism

In this work, we represent 3D body poses in terms of skeletons, such as those shown in Fig. 1, and the 3D locations of their D joints relative to that of a root node. As several authors before us [4, 16], we chose this representation because it is well adapted to regression and does not require us to know *a priori* the exact body proportions of our subjects. It suffers from not being orientation invariant but using temporal information provides enough evidence to overcome this difficulty.

Let $\mathbf{I}_i$ be the i-th image of a sequence containing a subject and $\mathbf{Y}_i \in \mathbb{R}^{3\times D}$ be a vector that encodes the corresponding 3D joint locations. Regression-based discriminative approaches to inferring $\mathbf{Y}_i$ typically involve a paramet-

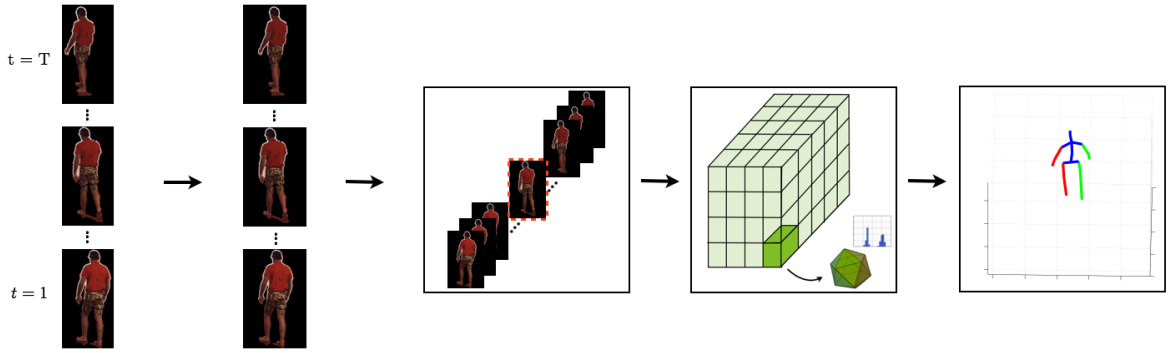

(a) Image stack (b) Motion compensation (c) Spatiotemporal block (d) Dense HOG3D (e) Regression to central 3D pose

Figure 2. Overview of our approach to 3D pose estimation. (a) A person is detected in several consecutive frames. (b) The corresponding image windows are shifted so that the subject remains centered. (c) A data volume is formed by concatenating these aligned windows. (d) A pyramid of 3D HOG features are extracted densely over the volume. (e) The 3D pose in the central frame is obtained by regression.

ric [1, 5, 17] or non-parametric [37, 23] model of the conditional distribution $p(\mathbf{Y}|\mathbf{X})$, where $\mathbf{X}_i = \Psi(\mathbf{I}_i; \mathbf{m}_i)$ is a feature vector computed over the bounding box or the foreground mask, $\mathbf{m}_i$, of the person in $\mathbf{I}_i$. The distribution parameters are usually learned from a labeled set of N training examples, $\mathcal{T} = \{(\mathbf{X}_i, \mathbf{Y}_i)\}_{i=1}^N$. As discussed in Section 2, in such a setting, reliably estimating the 3D pose is hard to do due to the inherent ambiguities of 3D human pose estimation such as self-occlusion and mirror ambiguity.

Instead, we model the posterior distribution conditioned on a data volume consisting of a sequence of T frames centered at image i, $\mathbf{V}_i = [\mathbf{I}_{i-T/2+1}, \mathbf{I}_{i-T/2+2}, \ldots, \mathbf{I}_{i+T/2}]$, that is, $p(\mathbf{Y}|\mathbf{Z})$ where $\mathbf{Z}_i = \xi(\mathbf{V}_i; \mathbf{m}_{i-T/2+1}, \mathbf{m}_{i-T/2+2}, \ldots, \mathbf{m}_{i+T/2})$ is a feature vector computed over the data volume. The training set becomes $\mathcal{T} = \{(\mathbf{Z}_i, \mathbf{Y}_i)\}_{i=1}^N$, where $\mathbf{Y}_i$ is the pose of the central frame in the image stack. In practice, we collect every block of consecutive T frames across all training videos to obtain data volumes, $\mathbf{V}_i$. We will show in the results section that this significantly improves performance and that the best results are obtained for sequences of $T = 24$ to 48 images, that is 0.5 to 1 second given the 50fps of the sequences of the Human3.6m [16] dataset.

3.2. Spatiotemporal Features

Our feature vector $\mathbf{Z}$ is based on the 3D HoG descriptor described in [39]. It is computed by first subdividing a data volume such as the one depicted by Fig. 2(c) into equally spaced bins. For each one, the histogram of oriented 3D spatio-temporal gradients [18] is then computed. To increase the descriptive power, we use a multi-scale approach. We compute several HoG descriptors using different bin sizes. In practice, we use 3 levels in the spatial dimensions—2×2, 4×4 and 8×8—and we set the temporal bin size to a small value—4 frames for 50 fps videos—to capture fine temporal details. Our final feature vector $\mathbf{Z}$ is obtained by concatenating these HoG descriptors into a single vector.

A strength of the 3D HoG descriptor is that it *simultaneously* encodes appearance and motion information in a relatively straightforward way. Not only does it preserve location and time dependent appearance information by concatenating the blocks into a single vector—as does the 2D HoG for appearance information only—but it also encodes motion information by computing temporal gradients. An alternative to encoding motion information in this way would have been to explicitly track body parts in the spatiotemporal volume, as is done in many of the methods we discussed in Section 2. However, tracking body parts in 2D is subject to ambiguities caused by the projection from 3D to 2D and we believe that not having to do this is a contributing factor to the good results we will show in Section 4. Furthermore, an interesting avenue for future research would be to explore if similar or even better results could be obtained by replacing HoG by potentially even more powerful CNNs operating on the spatiotemporal volume.

3.3. Alignment

For the 3D HoG descriptors introduced above to be representative of the person's pose, the temporal bins must correspond to specific body parts, which implies that the person should remain centered from frame to frame in the bounding boxes used to build the image volume. Therefore to go from detected bounding boxes, such as those depicted by Fig. 1(a) to a useful spatiotemporal volume such as the one of Fig. 1(c), we need to shift the bounding boxes as shown in Fig. 1(b). In Fig. 3, we illustrate this requirement by showing heat maps of the gradients across a sequence without and with motion compensation. Without it, the gradients are dispersed across the region of interest, which reduces the feature stability.

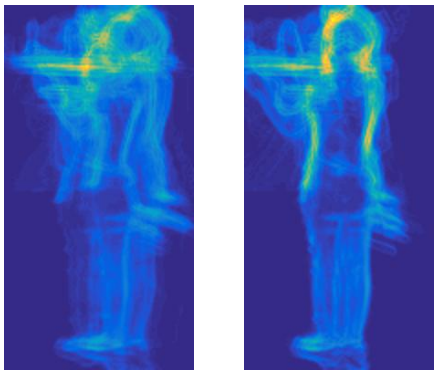

(a) No compensation (b) Motion compensation

Figure 3. Heat maps of the gradients across all frames for *Greeting* action (a) without motion compensation and (b) with motion compensation. Compensating for the motion, body parts become covariant with the HOG3D bins across frames and thus the extracted spatiotemporal features become more part-centric and stable.

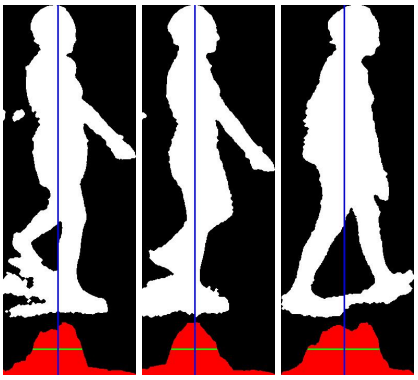

Figure 4. Motion compensation scheme. Histograms of foreground masks are computed for consecutive frames in a data volume and clipped with a threshold value. The mean of the maximum values in the clipped histogram is taken as the center of person. Clipping provides robustness against background subtraction artifacts and different articulations of the body. In order to have stability across frames, centers are temporally smoothed with a Kalman filter.

As will be discussed in more details in Section 4, when testing our approach on the HumanEva [31] and Human3.6m [16] datasets, we use the background subtraction masks or codes that they provide to find people's locations and position a bounding box around them. We resize all the images to the same size by keeping the aspect ratio of the person in it. Concatenating these bounding boxes does not guarantee a sufficiently good alignment across frames. To remedy this, we have explored the following two competing approaches.

1. We use a state-of-the-art optical flow-based motion stabilization algorithm [24].

2. When background subtraction is possible, we have implemented a simple scheme depicted by Fig. 4, which exploits the foreground mask produced by background subtraction. We first compute the histogram of pixel occurrences along each bounding box column. The histogram is then clipped with a cut-off value. Finally, we center the bounding box on the column that corresponds to the mean of the maximum values in the clipped histogram which are temporally smoothed with a Kalman filter.

In situations where background subtractions can be used, we will see that the second scheme, while simpler than the first, yields better results. However, the first one has the clear advantage that it does not depend on background subtraction and is therefore more widely applicable.

3.4. Regression

We cast 3D pose estimation in terms of finding a mapping $\mathbf{Z} \rightarrow \mathbf{f}(\mathbf{Z}) \approx \mathbf{Y}$, where $\mathbf{Z}$ is the descriptor computed over a spatiotemporal volume and $\mathbf{Y}$ is the 3D pose in its central frame. Following [16], we considered both an unstructured regression method—Kernel Ridge Regression [14], and a structured one—Kernel Dependency Estimation [8]—to learn f.

Kernel Ridge Regression (KRR) trains a model for each dimension of the pose vector separately. It first solves a regularized least-squares problem of the form

$$\underset{\mathbf{w}}{\operatorname{argmin}} \sum_i ||\mathbf{Y}_i - w_i k(\mathbf{Z}, \mathbf{Z}_i)||_2^2 + \lambda ||\mathbf{w}||_2^2 \ , \quad (1)$$

where the $(\mathbf{Z}_i, \mathbf{Y}_i)$ are training pairs, $\mathbf{w} = [w_1, \dots w_n]^\top$ are the parameters of the model, $k(\mathbf{Z}_1, \mathbf{Z}_2) = \exp(-\gamma \chi^2(\mathbf{Z}_1, \mathbf{Z}_2))$ is the exponential-χ^2 kernel [21], and λ is a regularization parameter. This can be done in closed-form by solving

$$\mathbf{w} = (\mathbf{K} + \lambda I)^{-1} \mathbf{Y} \ , \quad (2)$$

where $\mathbf{K}_{i,j} = k(\mathbf{Z}_i, \mathbf{Z}_j)$. Pose prediction is then carried out by computing for each $\mathbf{Z}$

$$\mathbf{f}(\mathbf{Z}) = \sum_i w_i \mathbf{K}(\mathbf{X}, \mathbf{X}_i) \ . \quad (3)$$

Kernel Dependency Estimation (KDE) is a structured regressor that accounts for correlations in 3D pose space, unlike conventional multiple output regression models, such as KRR. As argued in [16], it is more suitable to 3D human pose estimation due to the physical constraints and regularity of the human dynamics, which produce strong dependencies in joint positions.

To learn the regressor, the input and output vectors are first lifted into high-dimensional Hilbert spaces using kernel mappings Φ_Z and Φ_Y, respectively [8]. The dependency between high dimensional input and output spaces is modeled as a linear function. The corresponding matrix $\mathbf{W}$ is computed by standard kernel ridge regression, that is, by solving

$$\underset{\mathbf{W}}{\operatorname{argmin}} \sum_i ||\Phi_Y(\mathbf{Y}_i) - \mathbf{W}\Phi_Z(\mathbf{Z}_i)||_2^2 + \lambda||\mathbf{W}||_2^2 \ , \quad (4)$$

where λ is a regularization coefficient. To produce the final prediction $\mathbf{Y}$, the difference between the predictions and the mapping of the output in the high dimensional Hilbert space is minimized by finding

$$\underset{\mathbf{Y}}{\operatorname{argmin}} ||\mathbf{W}^T\Phi_Z(\mathbf{Z}) - \Phi_Y(\mathbf{Y})||_2^2 \ . \quad (5)$$

Although the problem is non-linear and non-convex, it can nevertheless be accurately solved given the KRR predictors for individual outputs to initialize the process.

In practice, we use an input kernel embedding based on 15,000-dimensional random feature maps corresponding to an exponentiated-χ^2 as in [21] and a 4000-dimensional output embedding corresponding to radial basis function kernel.

4. Results

In this section, we first introduce the datasets and parameters used in our experiments. Then, we describe the baselines we compare against and discuss our results.

4.1. Datasets

We evaluate our approach on two standard benchmarks for 3D human pose estimation. They are as follows:

Human3.6m is a recently released large-scale motion capture dataset that comprises 3.6 million images and corresponding 3D poses and complex motion scenarios. 11 subjects perform 15 different actions under 4 different viewpoints. The scenarios include a diverse set of motions of typical human activities, such as walking, eating, greeting, discussion etc.

HumanEva-I/II datasets provide synchronized images and motion capture data and are standard benchmarks for 3D human pose estimation. On these datasets, 3D joint positions are transformed to camera coordinates to provide monocular human pose estimation results.

For the experiments we carried out on these datasets, we use the average Euclidean distance between the ground truth and predicted joint positions as the evaluation metric. For the sake of completeness, we tabulate the parameters used in our experiments explained in Section 3 in Table 1.

Parameter:	Value
Spatial cell dividing	$2 \times 2, 4 \times 4, 8 \times 8$
Temporal cell size	4 frames
Temporal window size	24 frames
Input kernel dimensionality	15000
Output kernel dimensionality	4000

Table 1. The parameters used throughout our experiments.

4.2. Baselines

To demonstrate the effectiveness of our approach, we compare it against several state-of-the-art algorithms. We chose them to be representative of different approaches to 3D human pose estimation, as discussed in Section 2. For those which we do not have access to the code, we used the published performance numbers and ran our own method on the corresponding data. We list them below.

Single-frame based approaches rely on 2D appearance information for 3D human pose estimation. We will compare against the state-of-the-art 2D HoG based approach [16] and Convolutional Neural Networks [22] on Human3.6m dataset. We provide further comparisons for 3D Pictorial Structures [3] and Regression Forests [19] on the HumanEva-I dataset.

Frame-to-Frame Tracking Approaches use dynamical models to estimate transitions in consecutive frames and assume priors on the motion of the person. Among these approaches, we compare against Conditional Restricted Boltzmann Machines [36] and the loose-limbed body model of [33] that relies on probabilistic graphical models of the human pose and motion on HumanEva-I dataset.

Tracking-by-Detection Approaches refine single-frame detections by linking them into consistent object trajectories. Among these methods, we compare against the one of [2] on the HumanEva-II dataset.

4.3. Evaluation on the Human3.6m Dataset

To quantitatively evaluate the performance of our approach, we first used the recently released Human3.6m [16] large-scale motion capture dataset. On that dataset, the regression-based method of [16] performed best at the time and we therefore use it as a baseline. That method relies on a Fourier approximation of 2D HoG features using χ^2 comparison metric and we will refer to it as "$e^{\chi 2}$-HoG + KRR" or "$e^{\chi 2}$-HoG + KDE", depending on whether it uses the KRR or KDE regressors introduced in Section 3.4. Since then, even better results have been reported for some of the actions by using Convolutional Neural Nets (CNNs) [22].

Method:	Directions	Discussion	Eating	Greeting	Phone Talk	Posing	Buying	Sitting
e^{χ^2}-HoG + KRR [16]	140.00	189.36	157.20	167.65	173.72	159.25	214.83	193.81
e^{χ^2}-HoG + KDE [16]	132.71	183.55	132.37	164.39	162.12	150.61	171.31	151.57
CNN-Regression [22]	-	**148.79**	104.01	127.17	-	-	-	-
Ours + KRR	118.66	159.51	112.65	143.78	144.65	136.11	166.09	178.21
Ours + KDE	**102.39**	158.52	**87.95**	**126.83**	**118.37**	**114.69**	**107.61**	**136.15**
Method:	Sitting Down	Smoking	Taking Photo	Waiting	Walking	Walking Dog	Walking Pair	Average
e^{χ^2}-HoG + KRR [16]	279.07	169.59	211.31	174.27	108.37	192.26	139.76	178.03
e^{χ^2}-HoG + KDE [16]	243.03	162.14	205.94	170.69	96.60	177.13	127.88	162.14
CNN-Regression [22]	-	-	189.08	-	77.60	146.59	-	-
Ours + KRR	246.40	139.33	191.83	156.66	71.05	151.84	91.66	147.23
Ours + KDE	**205.65**	**118.21**	**185.02**	**146.66**	**65.86**	**128.11**	**77.21**	**125.28**

Table 2. 3D joint position errors using the metric of average euclidian distance between the ground truth and predicted joint positions to compare our results to those of two baselines [16, 22], when using different regressors as described in Section 3.4.

Given the current interest in Deep Learning approaches, we therefore use this method as a second baseline and we will refer to it as CNN-Regression.

The authors of [22] reported results on subjects S9 and S11, whereas those of [16] made their code available. To compare our results to those of both baselines, we therefore trained our regressors and those of [16] for 15 different actions. We used 5 subjects (S1, S5, S6, S7, S8) for training purposes and 2 (S9 and S11) for testing. The training and testing is carried out in all camera views for each separate action, as described in [16]. Recall from Section 3.1 that 3D body poses are represented by skeletons with 17 joints. Their 3D locations are expressed relative to that of a root node in the coordinate system of the camera that captured the images.

Table 2 summarizes our results on Human3.6m. Our method outperforms e^{χ^2}-HoG + KDE [16] significantly for all actions, with the mean error reduced by about 23%. It also improves on CNN-Regression [22] for the actions on which these authors reported accuracy numbers, except for "Discussion." Unsurprisingly, the improvement is particularly marked for actions, such as walking and eating, which involve substantial amounts of predictable motion. The discussion sequences, by contrast, involve motions that are very irregular and different from sequence to sequence as well as within sequences, which reduces the usefulness of temporal information. Example 3D pose reconstructions of our method for Human3.6m can be seen in Fig. 5. Further visualizations for the datasets are provided in the supplementary material.

Importance of Motion Compensation. To highlight the importance of motion compensation, we recomputed our features without it. As discussed in Section 3.3, we also tried two different approaches to performing it, using either a recent optical-flow based motion stabilization algorithm [24] or our own algorithm that exploits background subtraction results.

As shown in Table 3, motion compensation significantly improves performance for the two actions for which we performed this comparison. Furthermore, our approach to implementing it is more effective than the optical-flow method of [24].

Method:	*Greeting*	*Walking Dog*
e^{χ^2}-HoG [16]	164.39	177.13
Ours + No Compensation	144.48	138.66
Ours + Optical Flow [24]	140.97	134.98
Ours + Centering	**126.83**	**128.11**

Table 3. Importance of motion compensation. We compare the results of [16] against those obtained using our method, without motion compensation, with motion compensation using the algorithm of [24], and with motion compensation using the algorithm of Section 3.3. 3D joint position errors are given in millimeters.

Method:	*Walking*	*Eating*
e^{χ^2}-HoG + KDE [16]	96.60	132.37
Ours(12 frames)	69.11	94.32
Ours(24 frames)	65.86	87.95
Ours(36 frames)	**64.07**	87.64
Ours(48 frames)	64.75	**85.77**

Table 4. Influence of temporal window size. We compare the results of [16] against those obtained using our method with increasing temporal window sizes.

Influence of Temporal Window Size. In Table 4, we report the effect of changing the size of our temporal windows from 12 to 48 frames, also for two different actions. Since e^{χ^2}-HoG also relies on HoG features, its output can be thought of as the result for single-frame windows. Using temporal information clearly helps and the best results are

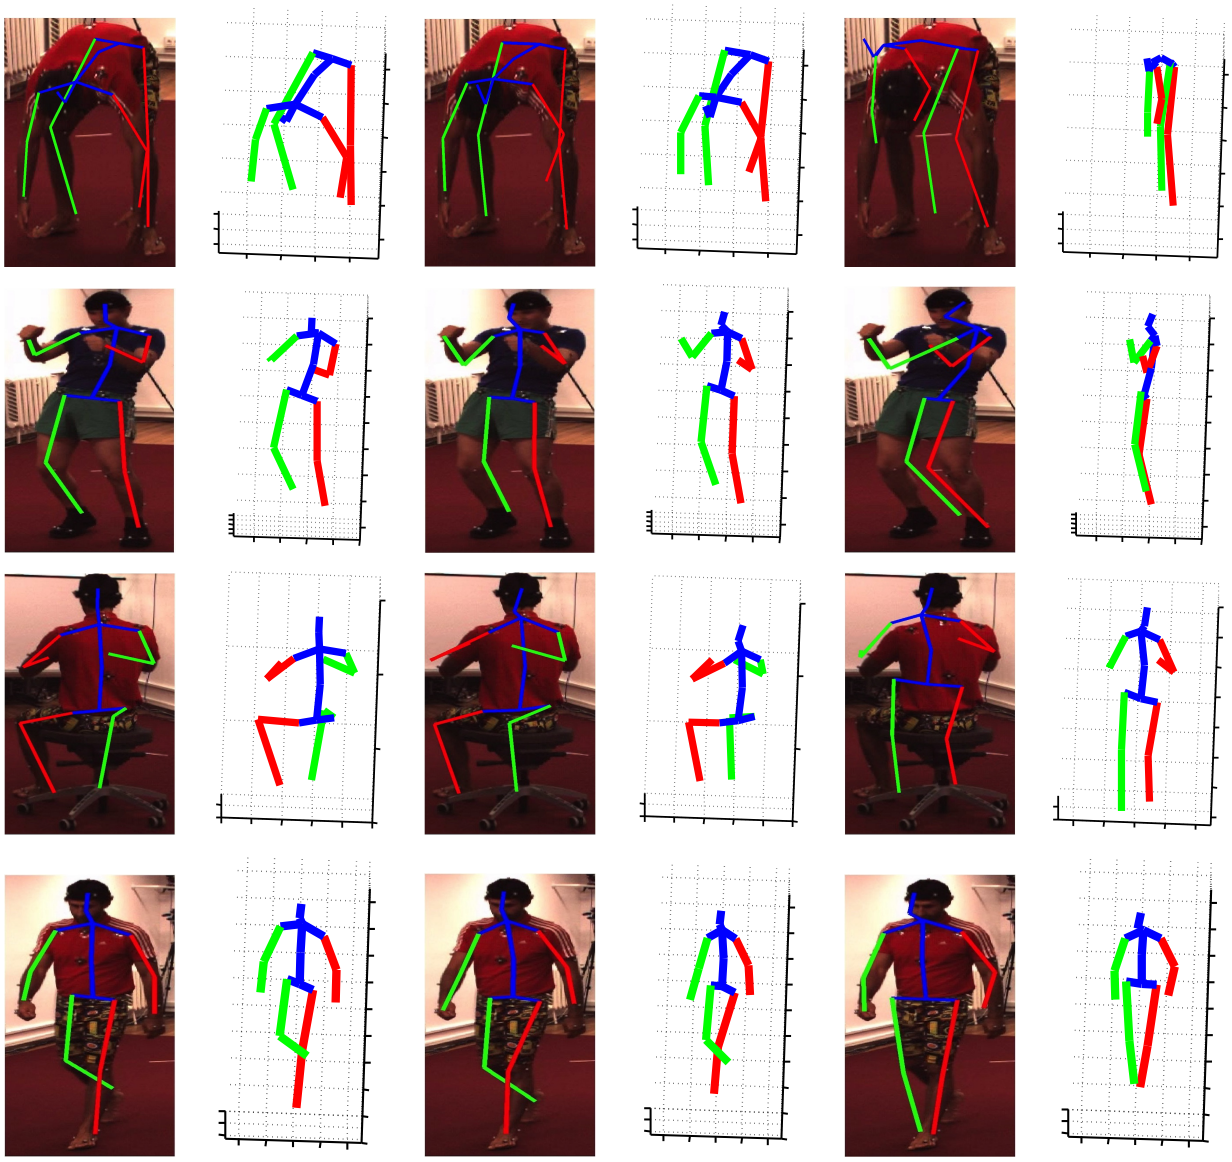

(a) Ground Truth (b) Our method (c) Ionescu et. al [16]

Figure 5. Example 3D pose estimation estimation results of challenging cases for *B*uying, *D*iscussion, *E*ating and *W*alking Pair actions in the Human3.6m dataset. The recovered 3D poses and their projection on the orthogonal plane is shown in **(b)** and **(c)** for our method and [16], respectively, along with the ground truth joint positions in **(a)**. Our method can recover the 3D pose of the person under these challenging scenarios where there is significant amount of self occlusion and orientation ambiguity, more accurately than [16]. Best viewed in color.

obtained in the range 24 to 48, which corresponds to 0.5 to 1 second at 50 fps. For the experiments we carried out on Human3.6m, we use 24 frames as it provides both accurate reconstructions and efficiency.

Activity Independent Regression. We further evaluate our approach in an experimental setting where all the motions from all the subjects are considered together for training. As for the activity-dependent setting, we use S1, S5, S6, S7, S8 for training and S9, S11 for testing. We obtain an average error of 121.62 mm across all test videos regardless of the action class, whereas the method of [16] yields an error of 154.43 mm. Interestingly, for both methods, this is better than the mean of the activity-dependent regressor as can be seen in the last column of the Table 2. This is due to the fact that the increase in the size of training samples compensate for the increase in the variance of 3D poses.

4.4. Evaluation on HumanEva Datasets

We further evaluated our approach on HumanEva-I and HumanEva-II datasets. Most of the early research in the field has reported results on this standard benchmark. Therefore we also carried out experiments on it to compare our approach to several state-of-the-art 3D human pose estimation and body tracking approaches. The baselines we considered include tracking-based approaches which impose dynamical priors on the motion [36, 33] and the tracking-by-detection framework of [2]. We demonstrate in Table 5 and Table 6 that using temporal information earlier in the inference process in a discriminative bottom-up fashion yields more accurate results than the above mentioned approaches that enforce top-down temporal priors on the motion.

HumanEva-I: For the experiments we carried out on HumanEva-I dataset, we train our regressor on training sequences of Subject 1, 2 and 3 and evaluate on the "validation" sequences in the same manner as the baselines we compare against [36, 33, 19, 3]. We report the performance of our approach on cyclic and acyclic motions, i.e., *W*alking and *B*oxing in Table 5. The results show that our proposed method outperforms state-of-the-art approaches.

	Walking			*Boxing*
Method:	S1	S2	S3	S1
Taylor et. al [36]	48.8	47.4	49.8	75.35
Sigal et. al [33]	66.0	69.0	-	-
Kostrikov et. al [19]	44.0	30.9	**41.7**	-
Belagiannis et. al [3]	68.3	-	-	62.70
Our	**32.6**	**24.1**	45.4	**46.18**

Table 5. Results on the *Walking* and *B*oxing sequences of the HumanEva-I dataset. We compare our approach against methods that rely on discriminative regression [19], 3D pictorial structures [3] and top-down temporal priors [36, 33].

HumanEva-II: On HumanEva-II, we compare against [2] as they have the best monocular pose estimation results on this dataset. Following the same experimental procedure, we use subjects $S1$, $S2$ and $S3$ from HumanEva-I dataset for training and report pose estimation results in the first 350 frames of the sequence containing subject S2. Whereas [2] uses additional training data from "People" [25] and "Buffy" [11] datasets, we only have to use the training data from HumanEva-I. We evaluated our approach using the official online evaluation tool. We illustrate the comparison in Table 6. For both camera views, our method is able to achieve significantly higher performance.

Method:	S2/C1	S2/C2
Andriluka et. al [2]	107	101
Our	**81.8**	**87.46**

Table 6. Results on the Combo sequence of the HumanEva-II dataset. We compare our approach against the tracking-by-detection framework of [2].

5. Conclusion

We have demonstrated that taking into account motion information very early in the modeling process yields substantial performance improvements over doing it *a posteriori* by linking detections in individual frames and only later imposing temporal consistency as part of the linking process.

This has been done using the hand-designed HOG 3D detector. However, the excellent performance of the method of [22] on ordinary 2D images suggests that even better results could be obtained by using instead Convolutional Neural Nets both for motion compensation and regression from the resulting spatiotemporal volumes to 3D poses. We will investigate this in future work.

References

[1] A. Agarwal and B. Triggs. 3D Human Pose from Silhouettes by Relevance Vector Regression. In *CVPR*, 2004.

[2] M. Andriluka, S. Roth, and B. Schiele. Monocular 3D Pose Estimation and Tracking by Detection. In *CVPR*, 2010.

[3] V. Belagiannis, S. Amin, M. Andriluka, B. Schiele, N. Navab, and S. Ilic. 3D Pictorial Structures for Multiple Human Pose Estimation. In *CVPR*, 2014.

[4] L. Bo and C. Sminchisescu. Twin Gaussian Processes for Structured Prediction. *IJCV*, 2010.

[5] L. Bo, C. Sminchisescu, A. Kanaujia, and D. Metaxas. Fast Algorithms for Large Scale Conditional 3D Prediction. In *CVPR*, June 2008.

[6] M. Burenius, J. Sullivan, and S. Carlsson. 3D Pictorial Structures for Multiple View Articulated Pose Estimation. In *CVPR*, 2013.

[7] X. Burgos-Artizzu, D. Hall, P. Perona, and P. Dollár. Merging Pose Estimates Across Space and Time. In *BMVC*, 2013.

[8] C. Cortes, M. Mohri, and J. Weston. A General Regression Technique for Learning Transductions. In *ICML*, 2005.

[9] J. Deutscher, A. Blake, and I. Reid. Articulated Body Motion Capture by Annealed Particle Filtering. In *CVPR*, 2000.

[10] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object Detection with Discriminatively Trained Part Based Models. *PAMI*, 32(9), 2010.

[11] V. Ferrari, M. Martin, and A. Zisserman. Progressive Search Space Reduction for Human Pose Estimation. In *CVPR*, 2008.

[12] J. Gall, B. Rosenhahn, T. Brox, and H.-P. Seidel. Optimization and Filtering for Human Motion Capture. *IJCV*, 2010.

[13] M. Hofmann and D. M. Gavrila. Multi-view 3D Human Pose Estimation in Complex Environment. *IJCV*, 2012.

[14] T. Hofmann, B. Schlkopf, and A. J. Smola. Kernel Methods in Machine Learning. *The Annals of Statistics*, 36(3):1171–1220, 2008.

[15] C. Ionescu, J. Carreira, and C. Sminchisescu. Iterated Second-Order Label Sensitive Pooling for 3D Human Pose Estimation. In *CVPR*, 2014.

[16] C. Ionescu, I. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large Scale Datasets and Predictive Methods for 3D Human Sensing in Natural Environments. *PAMI*, 2014.

[17] A. Kanaujia, C. Sminchisescu, and D. N. Metaxas. Semi-supervised Hierarchical Models for 3D Human Pose Reconstruction. In *CVPR*, 2007.

[18] A. Kläser, M. Marszałek, and C. Schmid. A Spatio-Temporal Descriptor Based on 3D-Gradients. In *BMVC*, 2008.

[19] I. Kostrikov and J. Gall. Depth Sweep Regression Forests for Estimating 3D Human Pose from Images. In *BMVC*, 2014.

[20] I. Laptev. On Space-Time Interest Points. *IJCV*, 64(2-3):107–123, 2005.

[21] F. Li, G. Lebanon, and C. Sminchisescu. Chebyshev Approximations to the Histogram χ^2 Kernel. In *CVPR*, 2012.

[22] S. Li and A. B. Chan. 3D Human Pose Estimation from Monocular Images with Deep Convolutional Network. In *ACCV*, 2014.

[23] R. Memisevic, L. Sigal, and D. J. Fleet. Shared Kernel Information Embedding for Discriminative Inference. *PAMI*, pages 778–790, April 2012.

[24] D. Park, C. L. Zitnick, D. Ramanan, and P. Dollár. Exploring Weak Stabilization for Motion Feature Extraction. In *CVPR*, 2013.

[25] D. Ramanan. Learning to Parse Images of Articulated Bodies. In *NIPS*, 2006.

[26] D. Ramanan, A. Forsyth, and A. Zisserman. Strike a Pose: Tracking People by Finding Stylized Poses. In *CVPR*, 2005.

[27] B. Sapp, A. Toshev, and B. Taskar. Cascaded Models for Articulated Pose Estimation. In *ECCV*, 2010.

[28] B. Sapp, D. J. Weiss, and B. Taskar. Parsing Human Motion with Stretchable Models. In *CVPR*, 2011.

[29] J. Shotton, A. Fitzgibbon, M. Cook, and A. Blake. Real-Time Human Pose Recognition in Parts from a Single Depth Image. In *CVPR*, 2011.

[30] L. Sigal, A. Balan, and M. J. Black. Combined Discriminative and Generative Articulated Pose and Non-rigid Shape Estimation. In *NIPS*, 2007.

[31] L. Sigal, A. Balan, and M. J. Black. Humaneva: Synchronized Video and Motion Capture Dataset and Baseline Algorithm for Evaluation of Articulated Human Motion. *IJCV*, 87(1-2):4–27, 2010.

[32] L. Sigal, S. Bhatia, S. Roth, M. J. Black, and M. Isard. Tracking Loose-limbed People. In *CVPR*, 2004.

[33] L. Sigal, M. Isard, H. W. Haussecker, and M. J. Black. Loose-limbed People: Estimating 3D Human Pose and Motion Using Non-parametric Belief Propagation. *IJCV*, 2012.

[34] C. Sminchisescu, A. Kanaujia, Z. Li, and D. Metaxas. Discriminative Density Propagation for 3D Human Motion Estimation. In *CVPR*, 2005.

[35] C. Sminchisescu and B. Triggs. Covariance Scaled Sampling for Monocular 3D Body Tracking. In *CVPR*, 2001.

[36] G. W. Taylor, L. Sigal, D. J. Fleet, and G. E. Hinton. Dynamical binary latent variable models for 3D human pose tracking. In *CVPR*, 2010.

[37] R. Urtasun and T. Darrell. Sparse Probabilistic Regression for Activity-Independent Human Pose Inference. In *CVPR*, 2008.

[38] R. Urtasun, D. Fleet, and P. Fua. 3D People Tracking with Gaussian Process Dynamical Models. In *CVPR*, 2006.

[39] D. Weinland, M. Ozuysal, and P. Fua. Making Action Recognition Robust to Occlusions and Viewpoint Changes. In *ECCV*, September 2010.

[40] F. Zhou and F. D. la Torre. Spatio-Temporal Matching for Human Detection in Video. In *ECCV*, 2014.

[41] S. Zuffi, J. Romero, C. Schmid, and M. J. Black. Estimating Human Pose with Flowing Puppets. In *ICCV*, 2013.