3D Human Body Tracking using Deterministic Temporal Motion Models

Raquel Urtasun and Pascal Fua [*]
Computer Vision Laboratory
Swiss Federal Institute of Technology
1015 Lausanne, Switzerland

Technical Report No: IC/2004/03

Abstract

There has been much effort invested in increasing the robustness of human body tracking by incorporating motion models. Most approaches are probabilistic in nature and seek to avoid becoming trapped into local minima by considering multiple hypotheses, which typically requires exponentially large amounts of computation as the number of degrees of freedom increases.

By contrast, in this paper, we use temporal motion models based on Principal Component Analysis to formulate the tracking problem as one of minimizing differentiable objective functions. The differential structure of these functions is rich enough to yield good convergence properties using a deterministic optimization scheme at a much reduced computational cost. Furthermore, by using a multi-activity database, we can partially overcome one of the major limitations of approaches that rely on motion models, namely the fact they are limited to one single type of motion.

We will demonstrate the effectiveness of the proposed approach by using it to fit full-body models to stereo data of people walking and running and whose quality is too low to yield satisfactory results without motion models.

1 Introduction

In recent years, much work has been devoted to increasing the robustness of people tracking algorithms by introducing motion models. Most approaches rely on probabilistic methods, such as the popular CONDENSATION algorithm [7, 5], to perform the tracking. While effective, such probabilistic approaches require exponentially large amounts of computation as the number of degrees of freedom in the model increases, and can easily become trapped into local minima unless great care is taken to avoid them [4, 3, 14, 15].

By contrast, in this paper, we use temporal motion models based on Principal Component Analysis (PCA) and inspired by those proposed in [12, 13] to formulate the tracking problem as one of minimizing differentiable objective functions. Our experiments show that the differential structure of these objective functions is rich enough to take advantage of standard *deterministic* optimization methods,

---

whose computational requirements are much smaller than those of *probabilistic* ones and can neverthe-less yield very good results even in difficult situations. Furthermore, in practice, we could combine both kinds of approaches [14].

We will further argue that we can partially overcome one of the major limitations of approaches that rely on motion-models, namely that they limit the algorithms to the particular class of motion from which the models have been created. This is achieved by performing PCA on motion databases that contain multiple classes of motions as opposed to a single one, which yields a decomposition in which the first few components can be used to classify the motion and can evolve during tracking to model the transition from one kind of motion to another.

We will demonstrate the effectiveness of the proposed approach by using it to fit full-body models to stereo data of people walking and running and whose quality is too low to yield satisfactory results without models. This stereo data simply provides us with a convenient way to show that this approach performs well on real data. However, any motion tracking algorithm that relies on minimizing an objec-tive function is amenable to the treatment we propose. We therefore view the contribution of this paper as the proposed formulation that produces results using a deterministic, as opposed to probabilistic optimization method, which yields good performance at a reduced computational cost.

In the remainder of this paper, we first discuss related approaches and our approach to body and motion modeling. We then introduce our deterministic optimization scheme and show its effectiveness using real data.

2   Related Work

Modeling the human body and its motion is attracting enormous interest in the Computer Vision com-munity, as attested by recent and lengthy surveys [6, 8]. However, existing techniques remain fairly brittle for many reasons: Humans have a complex articulated geometry overlaid with deformable tis-sues, skin and loosely-attached clothing. They move constantly, and their motion is often rapid, complex and self-occluding. Furthermore, the 3–D body pose is only partially recoverable from its projection in one single image. Reliable 3–D motion analysis therefore requires reliable tracking across frames, which is difficult because of the poor quality of image-data and frequent occlusions.

When a person is known *a priori* to be performing a given activity, such as walking or running, an effective means to constrain the search and increase robustness is to introduce a motion model. Of par-ticular interest to us are models that represent motion vectors as linear sums of principal components and have become widely accepted in the Computer Animation community as providing realistic results [1, 2]. These components are computed by capturing as many people as possible performing a specific activity, for example by means of an optical motion capture system, representing each motion as a temporally quantized vector of joint angles, and performing a Principal Component Analysis on the resulting set of vectors.

In practice, the position of a person, or body pose, in a given image frame can be defined by the position and orientation of a root node and a vector of joint angles. A motion can then be represented by an angular *motion vector*, that is a set of such joint angle vectors measured at regularly sampled intervals. Given a large enough database of motion vectors for different motion classes and the corresponding principal components $\Theta_j, 1 \leq j \leq m$, at a given time t, the joint angle vector $\Theta(\mu_t)$ can then be written

as

$$\Theta(\mu_t) = \Theta_0(\mu_t) + \sum_{j=1}^{m} \alpha_j \Theta_j(\mu_t) \quad \text{with } 0 \leq \mu_t \leq 1 \quad, \tag{1}$$

where μ_t is a normalized temporal variable that indicates to what stage of the motion the pose corresponds, Θ_0 represents an average motion, and the α_j are scalar coefficients. In short, the vector $(\mu_t, \alpha_1, ..., \alpha_m)$, where m is much smaller than the number of joint angles, can be used as the state vector that completely describes the body pose. Recovering this pose then amounts to minimizing an image-based objective function F with respect to this more compact representation, and can be expected to be much more robust than minimizing it with respect to the full set of joint angles.

This representation has already been successfully used in our community [12, 13], but always in a statistical context and without exploiting the fact that F is easily differentiable with respect to μ_t and the α_j coefficients of Eq. 1. Here, we propose to use this fact to formulate the fitting problem as a traditional optimization problem with respect to the $(\mu_t, \alpha_1, ..., \alpha_N)$ state vector. Instead of generating many "particles" by randomly choosing values for the α_j, we will compute the Jacobian of F and use it in conjunction with standard least-squares techniques [10].

3   Models

In this section, we introduce the models we use to describe both body pose and shape at a given time as well as its motion over time.

3.1   Body Model

In earlier work [9], we have developed a body-modeling framework that relies on attaching implicit surfaces, also known as soft objects, to an articulated skeleton. Each primitive defines a field function and the skin is taken to be a level set of the sum of these fields, as shown in Fig. 1a. Defining surfaces in this manner lets us define a distance function of data points to the model that is differentiable. We will take advantage of this to implement our minimization scheme, as discussed in Section 4.

As in Section 2, let us assume that, at a given time, the pose of the skeleton is entirely characterized by the global position and orientation G of a root node and a set of joint angles Θ. To avoid undue blending of primitives, the body is divided into several body parts. Each body part b includes n_b ellipsoidal primitives attached to the skeleton. To each primitive is associated a field function f_i of the form

$$f_i(G, \Theta, X) = b_i \exp(-a_i d_i(G, \Theta, X)) \quad, \tag{2}$$

where X is a 3–D point and d_i is the algebraic distance to this ellipsoid. The complete field function for body part b is taken to be

$$f^b(G, \Theta, X) = \sum_{i=1}^{n_b} f_i(d_i(X, G, \Theta)) \quad, \tag{3}$$

and the skin is the set

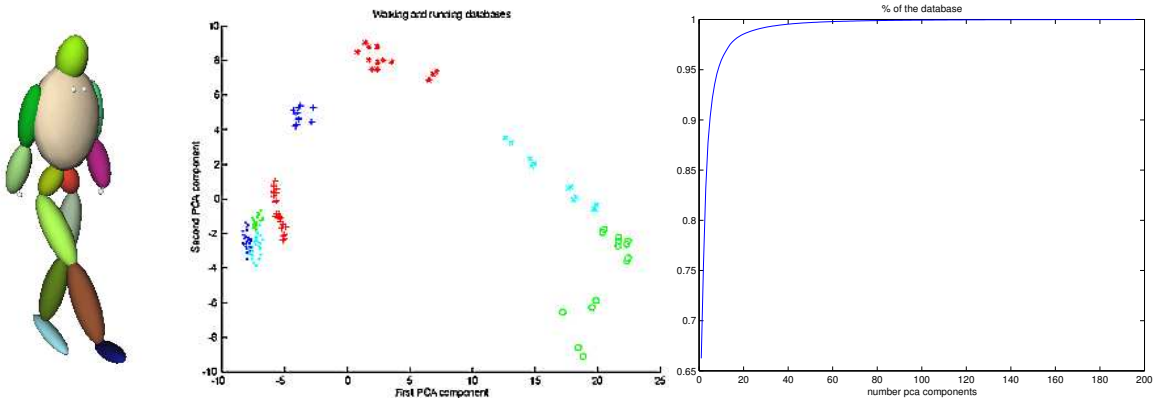$$\mathcal{S}(G, \Theta) = \bigcup_b \{X \in \Re^3 | f^b(G, \Theta, X) = C\} \quad, \tag{4}$$

3

Figure 1: Shape and motion models. a) Volumetric primitives attached to an articulated skeleton. b) First two PCA components for 4 different captures of 4 subjects walking at speeds varying from 3 to 7km/h, and running at speeds ranging from 6 to 12km/h. The data corresponding to different subjects is shown in different styles. c) Percentage of the database that can be generated with a given number of eigenvectors.

where C is a constant. A point X is said attached to body part b if

$$f^b(G, \Theta, X) = \min_{1 \leq i \leq B} |f^i(G, \Theta, X) - C| \tag{5}$$

Fitting the model to stereo-data acquired at time t then amounts to minimizing

$$F_t(G_t, \Theta_t) = \sum_{b=1}^{B} \sum_{X_t \in b} (f^b(G_t, \Theta_t, X_t) - C)^2 \quad , \tag{6}$$

where the X_t are the 3–D points derived from the data, each one being attached to one of the B body parts. Note that F_t is a differentiable function of the global position G_t and of the joint angles in Θ_t and that its derivatives can be computed fast [9].

3.2 Motion Models

To create a motion database, we used a Vicon^{tm} optical motion capture system and a treadmill to capture 4 people, 2 men and 2 women,

- walking at 9 different speeds ranging from 3 to 7 km/h, by increments of 0.5 km/h;

- running at 7 different speeds ranging from 6 to 12 km/h, by increments of 1.0 km/h.

The data was then segmented into cycles and normalized so that each one is represented by the same number of samples. To this end, spherical interpolation in quaternion space was used because it is the space in which a distance measuring the proximity of two orientations can be naturally defined. It therefore lets us interpolate with a meaningful angular velocity measure on an optimal path splining among orientation key frames [11]. Since people never perform the same motion twice in exactly the same fashion, we included in the database four walking or running cycles for each person and speed. The mean motion of the set of examples was subtracted and the M eigenvectors of Eq. 1 were obtained by

SVD. Retaining only $m \leq M$ eigenvectors, gave us a reduced base of the most significant subspace of the motion space, that is the one that contains σ % of the database, where λ_i is the i-th bigger eigenvalue.

$$\sigma = \frac{\sum_{i=1}^{m} \lambda_i}{\sum_{i=1}^{M} \lambda_i} \tag{7}$$

In our experiments, we chose $\sigma = 0.9$, which means that for the multi-activity database we need only 5 out of 256 coefficients, which corresponds to the total number of examples in the database. In Fig. 1c we display σ as a function of the number of eigenvectors. The same method was used for the walking and running databases independently. The estimation problem is thus reduced from the $\simeq 80$ degrees of freedom for the 28 joints in our body model at each time step, to 5 coefficients plus the time.

Fig. 1b shows the first two PCA components of the original examples used to create the joint walking and running database. The two activities produce separate clusters. The walking components appear on the left of the plot and form a relatively dense set. By contrast, running components are sparser because inter-subject variation is larger, indicating that more examples are required for a complete database.

Note that varying only the first two components along the curve corresponding to the path from one subset to another, yields very natural transitions between walking and running motions.

4 Deterministic Approach to Tracking

In this section we introduce our deterministic approach to tracking that relies on describing motion as a linear combination of the motion eigenvectors of Section 3.2 and choosing optimal weights for these vectors. As before, we represent the angular component of motion Θ as

$$\Theta = \Theta_0 + \sum_{i=1}^{m} \alpha_i \Theta_i \quad, \tag{8}$$

where Θ_0 is the average motion and the Θ_i are the eigenvectors of Section 3.2. Evaluating Θ at a particular time μ_t yields the pose

$$\Theta(\mu_t) = \Theta_0(\mu_t) + \sum_{i=1}^{m} \alpha_i \Theta_i(\mu_t) = [\theta^1(\mu_t), \cdots, \theta^{ndof}(\mu_t)]^T \quad, \tag{9}$$

where the θ^j are the actual joint angles at time μ_t for the $ndof$ degrees of freedom of the body model we use.

Note that the complete motion is described not only by the angular motion discussed above, but also by the motion G_t of the root body model node with respect to which the angles are expressed. This adds six degrees of freedom to our model, which are not represented at all in our motion database since the data was acquired on a treadmill on which the subjects were forced to walk straight. Furthermore, even if the global motion had been acquired, it would make no sense to include it in the database because similar motions would then have been considered as different just because of the orientation or position of the body.

Let us assume that we have acquired image data, which here is the stereo data depicted by Fig. 2 but could just as well be anything else, in T consecutive frames. Our goal is to recover the motion by minimizing an objective function F over all frames and, therefore, fitting the model to the image data.

Figure 2: Input stereo data. Top row: First image of a synchronized trinocular video sequence at three different times. The 3–D points computed by the Digiclops ^{tm} system are reprojected onto the images. Bottom row: Side views of these 3–D points. Note that they are very noisy and lack depth. A video is available in the supplementary material.

Tracking is achieved in two main steps. First the global motion G_t is recovered in a recursive way. Results from frame t are used as initialization for frame $t + 1$. We initialize using the average motion Θ_0, positioning the global motion for the first frame by hand, where the time $\mu_t, 1 \leq t \leq T$, is a linear interpolation between initial values μ_1 and μ_T for the first and last frames. For each frame we minimize

$$F_t(G_t, \Theta_0(\mu_t, \alpha_i)) \tag{10}$$

with respect to $G = G(t_x, t_y, t_z, \theta_x, \theta_y, \theta_z)$, where F_t is defined in equation 6. Given this global motion estimate, we then fit the data over all frames simultaneously by minimizing

$$F = \sum_{1 \leq t \leq T} F_t(G_t, \Theta(\mu_t, \alpha_i)) \ , \tag{11}$$

with respect to the μ_t, α_i and G_t.

In the remainder of this Section, for comparison purposes, we first show the result of fitting the stereo data we use without using motion models. We then introduce in more detail our approach to enforcing the motion models, with or without assuming that the style remains constant. Finally, we discuss the computational requirements of our scheme and contrast them with those of more traditional probabilistic approaches.

4.1 Tracking without Motion Models

In this paper we use stereo data acquired using a Digiclops^{tm} operating at a 640×480 resolution and a 14Hz framerate, which is relatively slow when it comes to capturing a running motion. The quality of the data is poor for several reasons. First, to avoid motion blur, we had to use a high shutter speed that

Figure 3: Tracking without a motion model. Given the low framerate, motion between frames is large enough to provoke erroneous attachments of data points to body parts and, as a consequence, very poor fitting behavior. The whole sequence is shown. An animated gif is available in the supplementary material.

reduces exposure too much. Second, because the camera is fixed and the subject must remain within the capture volume, she appears to be very small at the beginning of the sequence. As a result the data of Fig. 2 is very noisy and lacks both resolution and depth.

To establish a baseline, in Fig. 3, we show the unsatisfactory result of fitting our model to this data without using motion models, that is by minimizing the objective function of Eq. 6 in each frame separately. We simply use the recovered pose in frame $t - 1$ as the starting point in frame t.

Careful analysis shows that tracking fails chiefly due to the low framerate, as the interframe motion is too large. This prevents the process of "attaching" data points to body parts discussed in Section 3.1 from functioning properly. In the fifth image of Fig. 3, both legs end up being "attracted" to the same data points.

4.2   Tracking a Steady Motion

To remedy the problems discussed above, we can first assume that the motion is steady over T data frames and, therefore, that the α_i coefficients of Eq. 8 are invariant. The motion state vector is taken to be

$$\phi = [\overrightarrow{\mu}, \overrightarrow{\alpha}] = [\mu_1, \ldots, \mu_T, \alpha_1, \ldots, \alpha_m] \tag{12}$$

To effectively minimize the objective function F of Eq. 11 using a standard least-squares technique [10], we need to evaluate its Jacobian. Bearing in mind that the derivatives of F with respect to the individual joints angles $\frac{\partial F}{\partial \theta_j}$ can be easily computed [9], this can be readily done as follows:

$$\frac{\partial F}{\partial \alpha_i} = \sum_{j=1}^{ndof} \frac{\partial \theta_j}{\partial \alpha_i} \cdot \frac{\partial F}{\partial \theta_j} \ , \tag{13}$$

7

Figure 4: Motion vector and its temporal derivatives. Left: First 5 eigenvectors for the flexion-extension in the sagittal plane of the left knee. Right: Temporal derivatives $\frac{\partial \Theta_{ij}}{\partial \mu_t}$.

Figure 5: Tracking a walking motion assuming a constant style. The legs are correctly positioned. A video is available in the supplementary material.

$$\frac{\partial F}{\partial \mu_t} = \sum_{j=1}^{ndof} \frac{\partial \theta_j}{\partial \mu_t} \cdot \frac{\partial F}{\partial \theta_j} \ . \tag{14}$$

Because the θ_j are linear combinations of the Θ_i eigenvectors, $\frac{\partial \theta_j}{\partial \alpha_i}$ is simply the Θ_{ij}, the jth coordinate of Θ_i. Similarly, we can write

$$\frac{\partial \theta_j}{\partial \mu_t} = \sum_{i=1}^{m} \alpha_i \frac{\partial \Theta_{ij}}{\partial \mu_t} \ ,$$

where the $\frac{\partial \Theta_{ij}}{\partial \mu_t}$ can be evaluated using finite differences and stored when building the motion database, as shown in Fig. 4.

Fig. 5 displays the results on a walking sequence performed by a subject who was not recorded when building the database. To validate our results, he is wearing four gyroscopes on his legs, one for each sagittal rotation of the hip and knee joints. The angular speeds they measure are used solely for comparison purposes and we show their integrated values in Fig. 6. We overlay on the corresponding plots the values recovered by our tracker. Note that they are very close, even though the left leg is severely occluded.

Fig. 7 depicts results on the running sequence of Fig. 2 using the running database of Section 3.2, which are much better than those of Fig. 3. The pose of the legs is now correctly recovered, except the

Figure 6: Comparing recovered rotation angles using visual tracking, in blue, and integrating gyroscopic data, in red, for the walk of Fig. 5. Left column: Right hip and knee sagittal rotations. Right Column: Same thing for the left leg. Note that visual and gyroscopic plots are very close, even though the left leg is severely occluded most of the time.

Figure 7: Tracking a running motion assuming a constant style. The legs are correctly positioned except the left one in the first frame. Movies are submitted as supplementary material.

one of the left leg in the first frame. This is due in part to the fact that the database was acquired using a treadmill and is therefore too sparse to model a motion in which the leg is raised that high, and in part to the fact that the motion is not truly steady. We address these issues below.

4.3 Tracking a Variable Style and Speed Motion

In the sequences shown in this paper, speed and style are not truly constant. Because of space constraints, the subject starts, accelerates and stops over a short distance. This is true for walking and running, and even more so for transitions from one to the other. Using a single set of α_i parameters for the whole sequence as in Section 4.2 therefore overconstrains the problem. We relax these constraints by

9

Figure 8: Tracking a running motion while allowing the style to vary. The legs are now correctly positioned in the whole sequence. An animated gif is available in the supplementary material.

introducing a set of α_i per frame or per set of frames and the state vector then becomes:

$$\phi = \phi(\overrightarrow{\mu}, \overrightarrow{\alpha^1}, \ldots, \overrightarrow{\alpha^T}) \ \text{ where } \ \overrightarrow{\alpha^i} = (\alpha_1^i, \ldots, \alpha_m^i) \quad . \tag{15}$$

Improved tracking results from the running sequence of Fig. 2 are shown in Fig. 8. The system now has enough freedom to raise the leg in the first frame while still positioning the legs correctly everywhere else. Upper body tracking remains relatively imprecise because average errors in the stereo data are larger than the distance between torso and arms. Improving this would require the use of additional information, such as silhouette information, which could easily be done within the proposed framework. Similar results for walking are shown in Fig. 9. Small errors in foot positioning are due to the fact that ankle flexion has not been recorded in the motion database.

Having a set of PCA parameters per frame gives the system the freedom to automatically evolve from one activity to another. To demonstrate this, in Fig. 10, we use our full motion database to track a transition from walking to running. In the first few frames the subject is walking, then for a couple of frames she performs the transition and runs for the rest of the sequence. The arms are not tracked because we focus on estimating the motion parameters of the lower body only. Here again, the legs are successfully tracked with small errors in foot positioning that are due to the fact that ankle flexion is not part of the motion database.

4.4 Computational Requirements

Probabilistic approaches such as the one of [13] rely on randomly generating "particles" and evaluating their fitness. Assuming the cost of creating the particles to be negligible, the main cost of each iteration comes from evaluating an objective function, such as the function F of Eq. 11 for each particle. It is therefore in the order of $\mathcal{O}(npart(n))$ times the cost of computing F, where $npart$ is the number of particles, which grows very fast with n, the dimension of the state vector, to achieve good convergence properties.

Figure 9: Tracking a walking motion while allowing the style to vary. The sequence has a total of 18 frames, we show one in two. A video is available in the supplementary material.

By contrast, the main cost of each iteration of our optimization scheme comes from evaluating F and its Jacobian, which is of course more expensive than evaluating F alone. However, through careful implementation, we have found that it can be done at a cost in the order of $\mathcal{O}(\lg(ndof))$ times the cost of computing F alone, since evaluating F and its derivatives for the $ndof$ degrees of freedom in the body model involves many similar computations, and computing $\frac{\partial F}{\partial \theta_j}$ once per iteration is what is costly. As a consequence, the cost of the methods of Section 4.2 and 4.3 are of the same order and smaller than the probabilistic approach.

5 Conclusion

We have presented an approach using motion models that allows us to formulate the tracking problem as one of minimizing a differential objective function with respect to relatively few parameters. We take them to be the first few coefficients of the principal components of the joint angle space for motions captured using an optical motion capture device.

Using walking and running as examples, we have shown that this representation, while having a fairly low dimension, nevertheless has a rich enough differential structure to yield good performance at a low computational cost. It also has the ability to capture the transition from one motion to another.

Figure 10: Tracking the transition between walking and running. In the first four frames the subject is running. The transition occurs in the following three frames and the sequence ends with running. The whole sequence is shown. Corresponding movies are submitted as supplementary material.

Currently, the major limitation comes from the small size of the database we use, which we will endeavor to complete. This should allow us to precisely track a wider range of styles, perhaps at the cost of adding some regularization constraints that we presently do not need. We also plan to add additional motion types, such as jumping, for which motion capture data is fairly easy to acquire. In the current database, the samples corresponding to different people tend to cluster. If this remains true when the database is completed, this may become a promising approach not only for tracking but also for recognition.

References

[1] M. Alexa and W. Mueller. Representing animations by principal components. In *Eurographics*, volume 19, 2000.

[2] M.E. Brand and A. Hertzmann. Style Machines. *Computer Graphics, SIGGRAPH Proceedings*, pages 183–192, July 2000.

[3] K. Choo and D.J. Fleet. People tracking using hybrid monte carlo filtering. In *International Conference on Computer Vision*, Vancouver, Canada, July 2001.

[4] A. J. Davison, J. Deutscher, and I. D. Reid. Markerless motion capture of complex full-body movement for character animation. In *Eurographics Workshop on Computer Animation and Simulation*. Springer-Verlag LNCS, 2001.

[5] J. Deutscher, A. Blake, and I. Reid. Articulated Body Motion Capture by Annealed Particle Filtering. In *CVPR*, Hilton Head Island, SC, 2000.

[6] D.M. Gavrila. The Visual Analysis of Human Movement: A Survey. *Computer Vision and Image Understanding*, 73(1), January 1999.

[7] M. Isard. and A. Blake. CONDENSATION - conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1):5–28, August 1998.

[8] T.B. Moeslund and E. Granum. A Survey of Computer Vision-Based Human Motion Capture. *Computer Vision and Image Understanding*, 81(3), March 2001.

[9] R. Plänkers and P. Fua. Articulated Soft Objects for Multi-View Shape and Motion Capture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2003.

[10] W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes, the Art of Scientific Computing*. Cambridge U. Press, Cambridge, MA, 1992.

[11] K. Shoemake. Animating Rotation with Quaternion Curves. *Computer Graphics, SIGGRAPH Proceedings*, 19:245–254, 1985.

[12] H. Sidenbladh, M. J. Black, and D. J. Fleet. Stochastic tracking of 3D human figures using 2D image motion. In *European Conference on Computer Vision*, June 2000.

[13] H. Sidenbladh, M. J. Black, and L. Sigal. Implicit Probabilistic Models of Human Motion for Synthesis and Tracking. In *European Conference on Computer Vision*, Copenhagen, Denmark, May 2002.

[14] C. Sminchisescu and B. Triggs. Covariance Scaled Sampling for Monocular 3D Body Tracking. In *Conference on Computer Vision and Pattern Recognition*, Hawaii, 2001.

[15] C. Sminchisescu and B. Triggs. Kinematic Jump Processes for Monocular 3D Human Tracking. In *Conference on Computer Vision and Pattern Recognition*, Madison, WI, June 2003.