3D Head Trajectory Using a Single Camera

Caroline Rougier and Jean Meunier

Department of Computer Science and Operations Research (DIRO),
University of Montreal, QC, Canada
rougierc@iro.umontreal.ca, meunier@iro.umontreal.ca

Abstract. Video surveillance applications need tools to track people trajectories. We present here a new method to extract the 3D head trajectory of a person in a room using only one calibrated camera. A 3D ellipsoid representing the head is used to track the person with a hierarchical particle filter. This method can run in quasi-real-time providing reasonable 3D errors for a monocular system.

Keywords: computer vision, 3D, head tracking, monocular, particle filter, video surveillance.

1 Introduction

1.1 Motivation

A 3D trajectory gives a lot of information for behavior recognition and can be very useful for video surveillance applications. For instance, in our research on fall detection [17], the 3D head trajectory of a person in a room is needed. Indeed, Wu [11] shows in a biomechanical study with wearable markers that the 3D head velocities are efficient to detect falls. Thus, we need a method to extract the 3D head trajectory in real-time with sufficient precision to compute 3D velocities. In this paper, we present an original method which can track the person head with a 3D ellipsoid using a single calibrated camera and recover its location in a room.

1.2 Related Works

Usually, a multi-cameras system is used to have some 3D information. Wu and Aghajan [7] have recovered both the 3D trajectory of the head and its pose with a multi-camera system without accurate calibration of the camera network. Their method was based on Chamfer matching refined with a probabilistic framework. Kawanaka et al. [4] have tracked the head in a 3D voxel space with four stereo cameras, using particle filtering with depth and color information. Kobayashiet et al. [5] have used AdaBoost-based cascaded classifiers to extract the 3D head trajectory using several cameras.

For a low-cost system, we prefer to use only one camera per room, which explains our interest in monocular methods. Very few attempts have been done to track the 3D head trajectory in real-time with a single camera. Hild [3] has

recovered a 3D motion trajectory from a walking person, but he has supposed that the person was standing and that the camera optical axis was parallel to the (horizontal) ground plane. These assumptions can not be assumed in video surveillance applications: the camera need to be placed higher in the room to avoid occluding objects and to have a larger field of view, and the person is not supposed to be standing or facing the camera.

1.3 Method Overview

Birchfield [1] showed that a head is well-approximated by an ellipse in the 2D image plane. Its tracker was based on a local search using gradient and color information. A head ellipse was also tracked by Nummiaro *et al.* [10] using a color-based particle filter. In our previous work [17], we also have used this idea to compute the 3D head localization from a 2D head ellipse tracked in the image plane. The 3D head pose was computed using a calibrated camera, a 3D model of the head and its projection in the image plane. This method worked well with a standing person, but the 3D pose can be wrongly estimated when the person falls (the ellipse ratio can change with the point of view).

In this work, we explore the inverse method: the head is seen as a 3D ellipsoid which is projected as an ellipse in the 2D image plane, and tracked through the video sequence using a particle filter. Section 2 describes the 3D head model and its projection in the image plane. The head tracking module using particle filter is shown in section 3 and is tested in section 4 with the HumaneEva-I dataset [14].

2 Head Model Projection

We chose to represent the head by a 3D ellipsoid, which is projected in the image plane as an ellipse [6]. The 3D head model will be tracked in the world coordinate system attached to the XY ground plane. To project the 3D head model in the image plane, we need to know the camera characteristics (intrinsic parameters) and the pose of the XY ground plane related to the camera (extrinsic parameters).

The intrinsic parameters can be computed using a chessboard calibration pattern and the camera calibration toolbox for Matlab [16]. The camera's intrinsic matrix K is defined by:

$$K = \begin{bmatrix} f_x & 0 & u_0 & 0; & 0 & f_y & v_0 & 0; & 0 & 0 & 1 & 0 \end{bmatrix} \quad (1)$$

with the focal length (f_x, f_y) and (u_0, v_0) in pixels. Notice that image distortion coefficients (radial and tangential distortions) are also computed and used to correct the images for distortion before processing with our methodology.

The extrinsic parameters are obtained from corresponding points between the real word and the projected image points. The plane-image homography obtained from these two sets of points is used to computed the extrinsic parameters $M_{ext} = [R \ T; 0 \ 0 \ 0 \ 1]$, where R and T are respectively a 3D rotation matrix and a 3D translation vector, as described in the work of Zhang [12].

An ellipsoid is a quadric described by a matrix Q_C in the camera coordinate system such that $[x, y, z, 1]^T Q_C [x, y, z, 1] = 0$ for a point (x, y, z) belonging to the ellipsoid. This ellipsoid is projected in the image plane with the projection matrix P as a conic C [6][13]:

$$C = Q_{C_{44}} Q_{C_{1:3,1:3}} - Q_{C_{1:3,4}} Q_{C_{1:3,4}}^T \quad (2)$$

The ellipse is described by the conic by $[u, v, 1]^T C [u, v, 1] = 0$ for a point (u, v) in the image plane. Concretely, the ellipsoid matrix representing the head in the head coordinate system has the form:

$$Q_H = \begin{pmatrix} \frac{1}{B^2} & 0 & 0 & 0 \\ 0 & \frac{1}{B^2} & 0 & 0 \\ 0 & 0 & \frac{1}{A^2} & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (3)$$

With the semi-major A and the semi-minor B ellipsoid head axes.

The head ellipsoid is projected as Q_C in the camera coordinate system with the projection matrix P :

$$Q_C = P^{-1T} Q_H P^{-1} \quad (4)$$

The projection matrix P represents here the transformation from the head ellipsoid coordinate system to the image plane:

$$P = K M_{ext} M_{Head/World} \quad (5)$$

$M_{Head/World}$ will be defined during the tracking process by the translation and the rotation of the head in the world coordinate system (see Section 3). Finally, the conic obtained from eq. 2 is used to defined the ellipse parameters. An example of ellipsoid projection is shown in Fig. 1.

3 3D Head Tracking with Particle Filter

3.1 Particle Filter

Particle filters have been used with success in several applications, and in particular to track a head with an ellipse [17][10], or a parametric spline curve [8] using color information or edge contours. In our work, a particle filter is particularly suitable as it allows abrupt variations of the trajectory and can deal with small occlusions.

Fig. 1. The 3D ellipsoid model (a) is projected into an ellipse in the image plane (b). From the foreground silhouette (c), the foreground coefficient is computed using the normal segments to the ellipse. For example, in this case, we obtain a foreground coefficient equal to $C_F = 0.81$ (see Sect. 3.2).

The main idea of particle filters is to estimate the probability distribution $p(X_t|Z_t)$ of the state vector X_t of the tracked object given Z_t , representing all the observations. This probability can be approximated from a set of N weighted samples $S_t = \{s_t^{(n)}, \pi_t^{(n)}, 1 \leq n \leq N\}$ (also called particles). The main steps of the particle filter algorithm are:

Selection. N new samples are selected from the old sample set by favoring the best particles.

Prediction. The new samples are predicted with a stochastic dynamical model $S_t'^{(n)} = AS_t^{(n)} + Bw_t^{(n)}$ where $w_t^{(n)}$ is a vector of standard normal random variables, A and B are, respectively, the deterministic and stochastic components of the dynamical model.

Measurement. A weight $\pi_t^{(n)}$ is assigned to each sample (see Sect. 3.2).

Estimation. The mean state is then estimated using the weighted particles.

In our work, each particle of the filter is an ellipsoid, represented by the state vector:

$$X = [X_e, Y_e, Z_e, \theta_{X_e}, \theta_{Y_e}] \quad (6)$$

where (X_e, Y_e, Z_e) is the 3D head ellipsoid centroid expressed in the world coordinate system (translation component of the matrix $M_{Head/World}$), and $(\theta_{X_e}, \theta_{Y_e})$ are respectively the rotation around the X and the Y axes (rotation component of the matrix $M_{Head/World}$)¹.

¹ Notice that two angles (instead of three) are sufficient to define the position and orientation of the ellipsoid since its minor axes have both the same length in our model.

3.2 Particles Weights

The particle weights are based on foreground, color and body coefficients.

- **Foreground coefficient C_F**

To compute the foreground coefficient, we first need to extract the person silhouette in the image. For this purpose, we use a background subtraction method which consists in comparing the current image with an updated background image. We use the codebook method from the work of Kim *et al.* [9] which takes into account shadows, highlights and high image compression. The foreground coefficient is computed by searching for silhouette contour points along N_e ellipse normals. For each normal, the distance d_e from the ellipse point to the detected silhouette point is used to compute the foreground coefficient:

$$C_F = \frac{1}{N_e} \sum_{n=1}^{N_e} (D_e(n) - d_e(n)) / D_e(n) \quad , \quad C_F \in [0 \dots 1] \quad (7)$$

Where D_e , the half length of the normal, is used to normalize the distances. An example of foreground silhouette obtained with the codebook method [9] is shown in Fig. 1. This figure shows also the foreground coefficient obtained for a detected ellipse.

- **Color coefficient C_C**

The color coefficient is based on the normalized 3D color histogram of the head ellipse [10]. The color histogram, in the RGB color space, is composed of $N_b = 8 \times 8 \times 8$ bins and is computed inside a rectangular zone included in the ellipse. The comparison between an updated color head model and the target model is then done by calculating the histogram intersection:

$$C_C = \sum_{i=1}^{N_b} \min(H(i), H_{ref}(i)) \quad , \quad C_C \in [0 \dots 1] \quad (8)$$

- **Body coefficient C_B**

The body coefficient is used to link the head to the body through the body center. The projection of the 3D point corresponding to the centroid of the person should be near the centroid of the 2D silhouette (distance d_b compared to the half-major axis of the bounding box D_b). This coefficient is only used when the bounding box is valid (not used in case of occlusion for example).

$$C_B = (D_b - d_b) / D_b \quad , \quad C_B \in [0 \dots 1] \quad (9)$$

- **Final coefficient**

Finally, the ellipsoid coefficient is a combination of these coefficients amplified by a Gaussian to give larger weights to the best particles:

$$C_{final} = \frac{1}{\sqrt{2\pi}\sigma} \exp^{(C_F C_C C_B) / 2\sigma^2} \quad \text{with } \sigma = 0.25 \quad (10)$$

3.3 Tracking

Initially, a head detection module is used to calibrate our system. When the person stands up, the top head point is detected in the foreground silhouette. From this point, several 2D ellipses are tested and the one which has the biggest foreground coefficient C_F is kept. If $C_F > 0.7$, the ellipse is supposed sufficiently reliable. This head ellipse is used to calibrate the ellipsoid proportion related to the human height (The height H of the person is supposed to be known and the aspect ratio of the ellipse is fixed at 1.2 for a human head [1]). Finally, to obtain the initial 3D position of the ellipsoid, we use a floor-image homographic transformation of the image position of the feet to get the (X, Y) location and the person height H to infer the Z coordinate.

To have a reliable 3D head localization, we need to precisely estimate the head projection in the image. With a conventional particle filter, a lot of particles are needed for precision, which severely affects the computational performance and is incompatible with real-time operation. We rather prefer using an improved particle filter based on a hierarchical scheme with several layers, similar to the annealed particle filter [2]. We chose to work with 250 particles and 4 layers, which is a good compromise between performance and computational time.

The known velocity between two successive images is added to the particles to predict the next 3D ellipsoid localization before propagating the particles. Each layer has a different stochastic component $B = [B_{X_e}, B_{Y_e}, B_{Z_e}, B_{\theta_{X_e}}, B_{\theta_{Y_e}}]$ for the model propagation, sufficiently large for the first layer and decreasing for the next layers, such as $B_{l+1} = B_l/2$ with l the current layer and $l+1$ the next layer.

At the beginning, to refine the initial head position, B will be large for the X and Y components (initially, the person is supposed to be standing up, so that Z_e is approximately known and, θ_{X_e} and θ_{Y_e} are close to zero). For the next images, B is reinitialized from the current velocity between two successive images and fixed to $0.1m$ or $0.1rad$ if the velocity is too small.

4 Experimental Results

The precision of our method is evaluated with the HumanEva-I dataset [14], which contains synchronized multi-view video sequences and corresponding MoCap data (ground truth 3D localizations). As our method works with a single camera, we compared the results obtained from three different view points (color cameras C_1 , C_2 and C_3) using the video sequences of 3 subjects (S1, S2 and S3). The HumanEva dataset contains several actions. We used the motion sequences “walking” and “jog” to evaluate our 3D trajectories at 30Hz, 20Hz and 10Hz.

Figure 2 gives an example of 3D head trajectories obtained from different view points, and shows that the curves are similar to the MoCap trajectory. Table 1 summarizes the 3D mean errors obtained for each subject and each camera. This corresponds to a mean error of about 5% at a 4 to 6 meters distance. As expected, when the movement was larger, the error tended to be slightly higher, but the 3D ellipsoid remained well tracked.

Fig. 2. 3D head trajectories for a walking sequence of subject S1 (20Hz)

Table 1. Mean 3D errors (in cm) obtained from walking and jog sequences for different subjects (S1, S2, S3) and several view points (C1, C2, C3)

Frame rate	Camera	Action “walking”			Action “jog”		
		S1	S2	S3	S1	S2	S3
30Hz	C1	23.4 ± 15.6	23.2 ± 9.2	23.6 ± 12.6	17.8 ± 9	25.7 ± 18.2	22.3 ± 16
	C2	17.5 ± 14.6	24 ± 9.5	21.4 ± 11	27 ± 11.6	26 ± 9	17 ± 9
	C3	21 ± 13.7	24.1 ± 9.8	22.9 ± 13.6	27 ± 12.1	21.5 ± 9.4	33.1 ± 21.7
20Hz	C1	23.4 ± 14.2	22.9 ± 9.3	24.1 ± 11.9	21.4 ± 14.5	28.8 ± 22.3	26.4 ± 21.1
	C2	20.7 ± 11.4	22.9 ± 8.4	19.5 ± 9	30 ± 11.6	29.5 ± 12.4	18.8 ± 13.6
	C3	19.7 ± 9.8	23.1 ± 10.8	24.6 ± 13.1	23.4 ± 8.7	25.8 ± 12.1	30.9 ± 17.6
10Hz	C1	22.6 ± 12.3	24.1 ± 12.8	21.2 ± 15.6	20.7 ± 16.8	30.8 ± 24.4	51.3 ± 45.7
	C2	33.3 ± 15.9	23.1 ± 7.7	21.2 ± 15.3	35.6 ± 15.4	34.4 ± 15.4	37.8 ± 23.5
	C3	21.7 ± 14	20.4 ± 9.5	25.7 ± 12.7	18.5 ± 8.2	35.3 ± 12.5	39.2 ± 28.9

Our method is implemented in C++ using the OpenCV library [15] and can run in quasi-real time (130ms/frame on an Intel Core 2 Duo processor (2.4 GHz), not optimized code and image size of 640x480).

5 Discussion and Conclusion

In this paper, we have shown that a 3D head trajectory can be computed with only one single calibrated camera. The method gave similar results for different viewpoints, different frame rates and different subjects. The 3D locations were estimated with a reasonable mean error of around 25cm (5% at 5 meters) which is sufficient for most activity recognition based on trajectories.

Acknowledgement

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

1. Birchfield, S.: Elliptical head tracking using intensity gradients and color histograms. In: Proc. IEEE Computer Vision and Pattern Recognition (CVPR), pp. 232–237 (1998)
2. Deutscher, J., Blake, A., Reid, I.D.: Articulated Body Motion Capture by Annealed Particle Filtering. In: Proc. IEEE Computer Vision and Pattern Recognition (CVPR), vol. 2, pp. 126–133 (2000)
3. Hild, M.: Estimation of 3D motion trajectory and velocity from monocular image sequences in the context of human gait recognition. In: International Conference on Pattern Recognition (ICPR), vol. 4, pp. 231–235 (2004)
4. Kawanaka, H., Fujiyoshi, H., Iwahori, Y.: Human Head Tracking in Three Dimensional Voxel Space. In: International Conference on Pattern Recognition (ICPR), vol. 3, pp. 826–829 (2006)
5. Kobayashi, Y., Sugimura, D., Hirasawa, K., Suzuki, N., Kage, H., Sato, Y., Sugimoto, A.: 3D Head Tracking using the Particle Filter with Cascaded Classifiers. In: Proc. of British Machine Vision Conference (BMVC), pp. 37–46 (2006)
6. Stenger, B., Mendonça, P.R.S., Cipolla, R.: Model-Based Hand Tracking Using an Unscented Kalman Filter. In: Proc. BMVC, vol. 1, pp. 63–72 (2001)
7. Wu, C., Aghajan, H.: Head pose and trajectory recovery in uncalibrated camera networks - Region of interest tracking in smart home applications. In: ACM/IEEE International Conference on Distributed Smart Cameras, pp. 1–7 (2008)
8. Isard, M., Blake, A.: Condensation - conditional density propagation for visual tracking. *International Journal of Computer Vision* 29(1), 5–28 (1998)
9. Kim, K., Chalidabhongse, T.H., Harwood, D., Davis, L.: Real-time foreground-background segmentation using codebook model. *Real-Time Imaging* 11(3), 172–185 (2005)
10. Nummiaro, K., Koller-Meier, E., Van Gool, L.: An adaptive color-based particle filter. *Image and Vision Computing* 21(1), 99–110 (2003)
11. Wu, G.: Distinguishing fall activities from normal activities by velocity characteristics. *Journal of Biomechanics* 33(11), 1497–1500 (2000)
12. Zhang, Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(11), 1330–1334 (2000)
13. Hartley, R.I., Zisserman, A.: *Multiple view geometry in computer vision*, 2nd edn. Cambridge University Press, Cambridge (2004)
14. Sigal, L., Black, M.J.: HumanEva: Synchronized video and motion capture dataset for evaluation of articulated human motion Brown University, Department of Computer Science, Providence, RI, CS-06-08 (2006)
15. OpenCV: Intel Open Source Computer Vision Library, <http://www.intel.com/technology/computing/opencv>
16. Bouguet, J.Y.: Camera calibration toolbox for Matlab, http://www.vision.caltech.edu/bouguetj/calib_doc/index.html
17. Rougier, C., Meunier, J., St-Arnaud, A., Rousseau, J.: Monocular 3D Head Tracking to Detect Falls of Elderly People. In: International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 6384–6387 (2006)