Bartosz Tworek[1], Alexandre Bernardino[2] and José Santos-Victor[2]
[1] Faculty of Electrical Eng., Automatics, Computer Sc. and Electronics,
AGH University of Technology and Science, Krakow, Poland
bartosz.tworek@gmail.com
[2] Institute for Systems and Robotics,
Instituto Superior Técnico, Lisboa, Portugal
{alex,jasv}@isr.ist.utl.pt

BEST CONTROL/INSTRUMENTATION PAPER

VISUAL HOMING OF ROBOT HEADS WITHOUT ABSOLUTE SENSORS

ABSTRACT

With the increasing miniaturization of robotic devices, many actuators lack absolute position sensing. In these cases the initial position of the joints is unknown at startup. In this paper we present a vision based method for the automatic homing of serial kinematic structures composed of rotational joints, and having a perspective camera on the end-effector. Examples of such systems are pan-tilt surveillance cameras and most kinds of humanoid robot heads. The method is based on producing motions with known amplitude in one of the joints of the kinematic chain to induce image motion in the camera. The analysis of the induced homography allows the computation of the unknown angle of the other joint. The method can be iterated on more axes to calibrate longer serial chains. It requires calibrated cameras, static objects while homing and short links in the kinematics structure (or, equivalently, far away objects). We have implemented and validated the method in a small humanoid robot head.

1. INTRODUCTION

Many off-the-shelf DC motors are equipped with incremental encoders as the main feedback sensor, lacking absolute position sensing. These types of motors are often used in the construction of robots and other automated devices. In these systems it is not possible to know the initial position of the joints at startup and a procedure is necessary to set the robot to a known state, denoted as home or zero position.

To address this problem, it is usual to equip the robot with limit switches, or homing switches, that detect when the axes are in particular angular positions. However, due to miniaturization constrains, it may not be possible to install such sensors in the robot. Another possibility is to drive the axes to a mechanical stop and monitor the motor current. When the current exceeds a certain value, then the motor has reached the mechanical limit, whose angle can be known a priori. However, this procedure adds a source of physical stress in the system and may damage the mechanical components in the long term.

Even when the above strategies are feasible, they require the careful placement of limit and home switches, and a precise measurement of mechanical limits. Additionally, when attaching the cameras to the end-effector, there are always some misalignments that may degrade the initial calibration procedure.

In this paper we propose a solution to this problem for certain types of kinematics structures having cameras at the last joint of the chain. We present a self-homing procedure for system start-up that does not require absolute sensors, neither home/limit switches, nor the need to drive the system to mechanical hard stops. Instead, it performs small prospective motions in the robot joints and observes the image motion induced in the camera. It is assumed that the scene is static and that axes almost intersect (or, equivalently, objects are distant enough from the camera with respect to the length of the kinematics links). In these circumstances the induced image motion only depends on the given motion and the angle between the camera's optical axis and the rotation axis. By iterating this procedure in the several robot axes, it is therefore possible to automatically determine the wake-up state of the system.

There are very few works addressing the problem of visual based homing. Sometimes visual homing denotes the process of driving a system to some known position in the environment (see for instance [9]). In our case we drive the robot to a known kinematic configuration rather that a known position in space. To the best of our knowledge, the only work related to ours is [10] where the home configuration of a robot arm is achieved using images taken from outside cameras. In our case the cameras are "inside" the robotic system being calibrated.

The types of kinematic structures we consider are very common both in surveillance cameras and in robot heads. We implement the method and present results in a small humanoid robot head, calibrating its eyes and neck. Notwithstanding, the principle can be easily extended for other serial kinematics structures.

This paper is organized as follows. In Section 2 we formulate the problem in terms of the system kinematics and a homography estimation procedure. Then, in Section 3 we present a methodology to estimate the particular homographies arising in this problem. Section 4 is devoted to the presentation of experimental results of the application of the proposed method to the automatic homing of a small robot head. Finally, Section 5 presents the conclusion of the work and directions for future developments.

2. PROBLEM FORMULATION

In this section we formulate our problem in terms of a homography estimation problem. A homography is a transformation that is able to explain the relationship between the points observed in an image before and after a rotation of the camera. From the homography it is often possible to recover the rotation angles. Therefore, we are going to analyze the homography arising from the prospective motions applied to the robot, as a function of the initial, unknown, joint angles.

Let us consider initially the tilt-pan kinematics structure presented in Fig. 1. A camera is attached first to a pan unit, and then the pan unit is attached to a tilt unit. A similar analysis can be made for other kinematics structures.

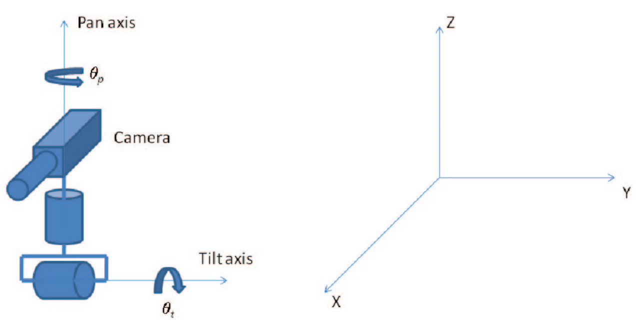

Figure 1 · Left: the kinematics structure of the pan-tilt system considered in this paper. Right: the adopted orientation for all the involved coordinate frames.

Considering identical reference coordinate frames for all joints in the canonical state, as shown in Fig. 1, the rotation matrix representing the camera's orientation with respect to the world reference frame depends on the pan and tilt angular displacements:

$$R_c(\theta_t,\theta_p) = Rot_y(\theta_t)Rot_z(\theta_p) =$$

$$= \begin{bmatrix} c_t & 0 & -s_t \\ 0 & 1 & 0 \\ s_t & 0 & c_t \end{bmatrix} \begin{bmatrix} c_p & -s_p & 0 \\ s_p & c_p & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} c_t c_p & -c_t s_p & -s_t \\ s_p & c_p & 0 \\ s_t c_p & -s_t s_p & c_t \end{bmatrix}$$

with $c_t = cos(\theta_t), s_t = sin(\theta_t),\ c_p = cos(\theta_p)$ and $s_p = sin(\theta_p)$.

Fixed points in the world, at coordinates *(X, Y, Z)* can be expressed in the camera frame by:

$$\begin{bmatrix} X_c(\theta_p,\theta_t) \\ Y_c(\theta_p,\theta_t) \\ Z_c(\theta_p,\theta_t) \end{bmatrix} = \left[R_c(\theta_p,\theta_t)\right]^T \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} c_t c_p X + s_p Y + s_t c_p Z \\ -c_t s_p X + c_p Y - s_t s_p Z \\ -s_t X + c_t Z \end{bmatrix}$$

The perspective projections of these points in the image plane have the following normalized coordinates:

$$x(\theta_p,\theta_t) = \frac{Y_c}{X_c} = \frac{-c_t s_p X + c_p Y - s_t s_p Z}{c_t c_p X + s_p Y + s_t c_p Z}$$

$$y(\theta_p,\theta_t) = \frac{Z_c}{X_c} = \frac{-s_t X + c_t Z}{c_t c_p X + s_p Y + s_t c_p Z}$$

Let us consider that, at start-up, the system has initial angles θ_t^0 and θ_p^0. These angles are unknown when the system is turned on. Then, a prospective motion of the tilt unit is performed: the tilt angle is changed by θ_t. This process is illustrated in Fig. 2.

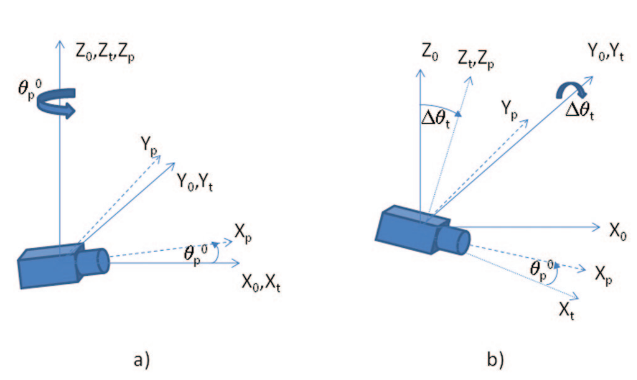

Figure 2 · The geometry of the system before (a) and after (b) the prospective motion.

For the sake of simplicity, and without loss of generality, we can consider a null initial tilt angle, $\theta_t^0 = 0$. This corresponds to set the world reference frame aligned with the initial robot's tilt frame. In the above conditions, points observed by the camera at system start-up are at coordinates:

$$x_0 = \frac{s_{p0}X + c_{p0}Y}{c_{p0}X + s_{p0}Y} \tag{1}$$

$$y_0 = \frac{Z}{c_{p0}X + s_{p0}Y} \tag{2}$$

After the tilt's prospective motion, these points are observed in new image coordinates:

$$x_1 = \frac{-c_t s_{p0}X + c_{p0}Y - s_t s_{p0}Z}{c_t c_{p0}X + s_{p0}Y + s_t c_{p0}Z} \tag{3}$$

$$y_1 = \frac{-s_t X + c_t Z}{c_t c_{p0}X + s_{p0}Y + s_t c_{p0}Z} \tag{4}$$

Let us recall that we are willing to estimate θ_{p0}, the unknown pan angle at start-up. θ_t is a known, actuated angle. The image coordinates x_0, y_0, x_1 and y_1 can be measured from the images by suitable image feature detectors and trackers, and *X*, *Y* and *Z* are unknown 3D coordinates of world points. Our objective now is to eliminate these coordinates in the previous equations.

From equations (1) and (2) we can write the following constraints:

$$\frac{Y}{X} = \frac{x_0 c_{p0} + s_{p0}}{c_{p0} - s_{p0}x_0} \tag{5}$$

$$\frac{Z}{X} = \frac{y_0}{c_{p0} - s_{p0}x_0} \tag{6}$$

Now, dividing both the numerator and denominator of Eqs. (3) and (4) by *X*, and introducing the constraints in Eqs. (5) and (6), we obtain:

$$x_1 = \frac{-c_t s_{p0}c_{p0} + c_t s_{p0}^2 x_0 + c_{p0}^2 x_0 + c_{p0}s_{p0} - s_t s_{p0}y_0}{c_t c_{p0}^2 - c_t c_{p0}s_{p0}x_0 + c_{p0}s_{p0}x_0 + s_{p0}^2 + s_t c_{p0}y_0} \tag{7}$$

$$y_1 = \frac{-s_t c_{p0} + s_t s_{p0}x_0 + c_t y_0}{c_t c_{p0}^2 - c_t c_{p0}s_{p0}x_0 + c_{p0}s_{p0}x_0 + s_{p0}^2 + s_t c_{p0}y_0} \tag{8}$$

These equations can be rewritten in the homography form:

$$[x_1 \quad y_1 \quad \lambda_1]^T = H[x_0 \quad y_0 \quad 1]^T$$

$$H = \begin{bmatrix} c_t s_{p0}^2 + c_{p0}^2 & -s_t s_{p0} & c_{p0}s_{p0}(1-c_t) \\ s_t s_{p0} & c_t & -s_t c_{p0} \\ c_{p0}s_{p0}(1-c_t) & s_t c_{p0} & c_t c_{p0}^2 + s_{p0}^2 \end{bmatrix} \tag{9}$$

A close inspection to the homography matrix shows that it has some repeated entries and only 6 of them are different. It has the form:

$$H = \begin{bmatrix} h_1 & -h_2 & h_3 \\ h_2 & h_4 & -h_5 \\ h_3 & h_5 & h_6 \end{bmatrix} \tag{10}$$

In the following section we will describe a method to estimate the entries of this matrix from the visual data. Once the homography is estimated, we can compute the unknown angle θ_{p0} by, e.g.:

$$\theta_{p0} = atan2(h_2,h_5)$$

3. COMPUTING THE HOMOGRAPHY

In this section we describe the method employed to estimate the particular homography arising in our formulation. The section is divided in three parts. The first part describes the methods employed to obtain point matches between the images before and after the prospective motion. The second part formulates the problem of estimating the homography from point matches, measured in the images. The third part presents the whole algorithm integrated in a robust estimation architecture (RANSAC).

3.1. Feature Tracker

For the estimation of the homography we need first to extract from the images a set of points visible on both images (before and after the prospective motion) and their pairwise correspondences:

$$\{(x_0^i, y_0^i) \mapsto (x_1^i, y_1^i), i = 1 \dots N\}$$

In the above equation, the lower index represents the image (0 for the image before and 1 for the image after the rotation), and the upper index represents the index of the point in the set, from 1 to N.

To obtain an adequate set of points in the first image, we use the corner detector in OpenCV [4], a open-source computer vision C library. The corner detector selects points in the image that are easy to track, and is based in [7]. It calculates the minimum eigenvalue of the Hessian matrix for each pixel of the image and then performs non-maxima suppression so that only local maxima in 3x3 neighborhood remains. Then it rejects the corners with the minimal eigenvalue less than a quality level determined by us. Finally it checks if all the corners found are separated enough one from another according to minimal distance determined by us.

The next step is to track the selected points. To do this we use the sparse iterative version of Lucas-Kanade optical flow method in pyramids [8], also implemented in the OpenCV library. It calculates the coordinates of the feature points on the current video frame given their coordinates on the previous frame. The function finds the coordinates with sub-pixel accuracy and rejects the points that cannot be reliably tracked in the second image.

3.2. Homography Estimation

There are several off-the-shelf routines for estimating homographies from point matches [5]. However, the homography arising in our problem has a special structure and we want to exploit this structure to improve its estimation. Using Eqs. (9) and (10) for each point match obtained in the tracking procedure, we have:

$$x_1^i = \frac{h_1 x_0^i - h_2 y_0^i + h_3}{h_3 x_0^i + h_5 y_0^i + h_6}, \quad y_1^i = \frac{h_2 x_0^i - h_4 y_0^i - h_5}{h_3 x_0^i + h_5 y_0^i + h_6}$$

Rearranging the previous equations, for each point we get the following constraints:

$$\begin{cases} -x_0^i h_1 + y_0^i h_2 + x_0^i x_1^i h_3 - h_3 + y_0^i x_1^i h_5 + x_1^i h_6 = 0 \\ -x_0^i h_2 + x_0^i y_1^i h_3 + y_0^i h_4 + y_0^i y_1^i h_5 + h_5 + y_1^i h_6 = 0 \end{cases}$$

This can be written vector form

$$\begin{cases} a_x^i h = 0 \\ a_y^i h = 0 \end{cases}$$

with

$$h = [h_1 \quad h_2 \quad h_3 \quad h_4 \quad h_5 \quad h_6]^T$$

$$a_x^i = [-x_0^i \quad y_0^i \quad x_0^i x_1^i - 1 \quad 0 \quad y_0^i x_1^i \quad x_1^i]$$

$$a_y^i = [0 \quad -x_0^i \quad x_0^i y_1^i \quad y_0^i \quad y_0^i y_1^i + 1 \quad y_1^i]$$

Given a set of N corresponding points, we can form the following linear system of equations:

$$Ah = 0$$

with

$$A = \begin{bmatrix} a_x^1 \\ a_y^1 \\ \vdots \\ a_x^N \\ a_y^N \end{bmatrix}$$

Since the homography has 6 different entries we need at least 3 points to estimate it (each point contributes with two equations). We can compute *h* through the Singular Value Decomposition (SVD) of A. From the SVD we take right singular vector which corresponds to the smallest singular value. Finally we can reshape the entries of h into the homography matrix *H*.

3.3. Robust Estimation

The above homography estimation method works well when there are no erroneous correspondences between the points in both images. Unfortunately, the tracking method sometimes provides false point matches that will degrade the results of the homography estimation. In order to address this problem, we use a well known robust estimation method that is able to eliminate the false matches (outliers) from the estimation process. The RANdom SAmple Consensus (RANSAC) [6] is an algorithm for robust estimation of models in the presence of many data outliers.

We apply the RANSAC algorithm to our problem as follows:

- Repeat for *L* times:
  - Select randomly a set of 3 feature pairs
  - Compute homography *H* using the selected feature pairs
  - Calculate a degree of fit (*d*) for all correspondences
  - Compute the number *K* of inliers consistent with *H*, i.e. the number of correspondences for which the degree of fit is lower than a threshold (*d*<*t*).
- Keep *H* whose set of inliers is the largest.

In our case, we use the symmetric transfer error to compute the degree of fit:

$$d = \sqrt{\|x - H^{-1}x'\|^2 + \|x' - Hx\|^2}$$

where $x \leftrightarrow x'$ are the point correspondences.

The number of iterations *L* is set adaptively, according to the following algorithm:

- $L = \infty$, *samplecounter* = 0
- While L = *samplecounter* repeat
  - Execute a RANSAC cycle (select random samples, compute homography and the number of inliers)
  - Estimate the proportion of outliers by $e = 1 - \frac{\#inliers}{\#totalpoints}$

— Set $L = \frac{log(1-p)}{log(1-(1-\varepsilon^{s}))}$ with $p = 0.99$ and $s = 3$ (number of feature pairs in a sample).

— Increment sample counter

• Terminate

This rule ensures that, with probability p, one among L samples will be free of outliers [5].

4. EXPERIMENTAL RESULTS

The proposed method was written in C++ and implemented as a module in the YARP framework [2]. YARP (Yet Another Robot Platform) is a middleware that facilitates the distributed processing and communication among different computers and provides operating system independence. We have also used the following additional libraries for the image based measurements and homography estimation:

- OpenCV (Open Source Computer Vision) which is used to create image processing part of the project [4].
- GSL (Gnu Scientific Library) [3] which is used to estimate the homography.

All used libraries are free open source software.

The method was then tested with the iCub humanoid robot's head [1]. It contains 6 DOFs: neck pan, tilt and swing and eye pan and tilt as shown in Fig. 3.

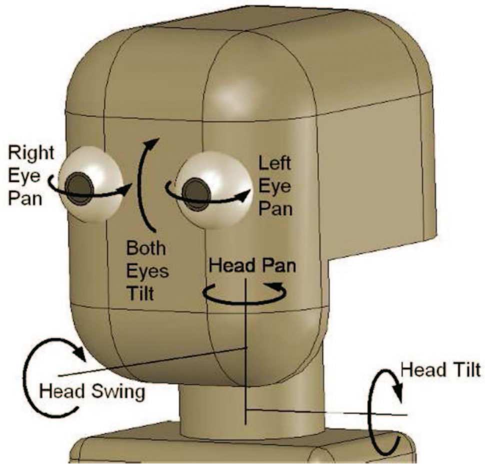

Figure 3 · Kinematics of the iCub robot head used in the tests.

We have performed the calibration of the right eye and head pan, which follow directly the formulation described in this paper. Whereas for calibrating the left eye the procedure is identical, for calibrating the other joints, a different set of equation is required, but can be derived in a straightforward manner using the same principles. The testing procedure was the following:

1. Find points of interest to track,
2. Move head or eye tilt,
3. Track points,
4. Compute homography matrix for correspondences,
5. Compute the pan angle.

Fig. 4 shows the sequence of actions required in the calibration procedure. Notice that the right eye, at startup, was offset from its canonical position. Then we have applied a tilt motion of 5 degrees. After the above steps, if we have a sufficient number of point matches (3 is the minimum) we are able to compute the homography and the initial pan angle. Fig. 5 shows the images and the tracked points used to compute the homography. Then we move the pan axis with the symmetrical of the computed value in order to set the eye to its canonical position (zero pan).

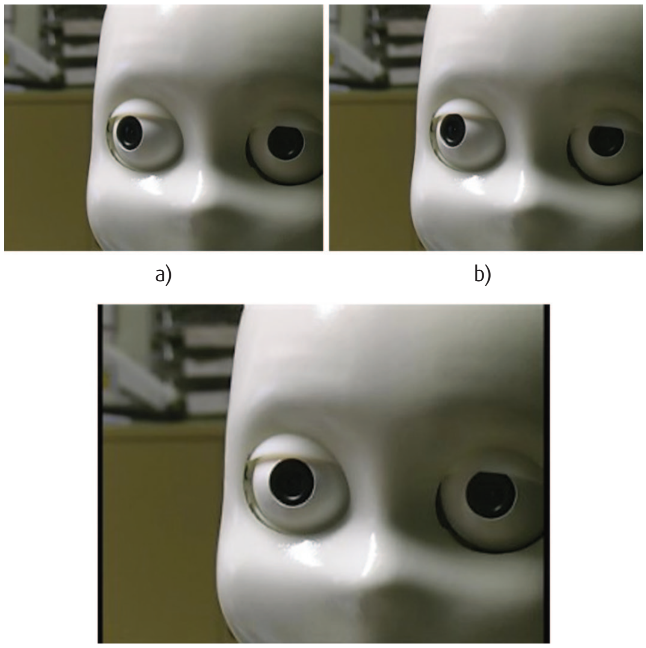

a) b)

c)

Figure 4 · Sequence of motions required to calibrate the eye. a) right eye position at startup. b) right eye position after prospective tilt motion. c) right eye position after calibration.

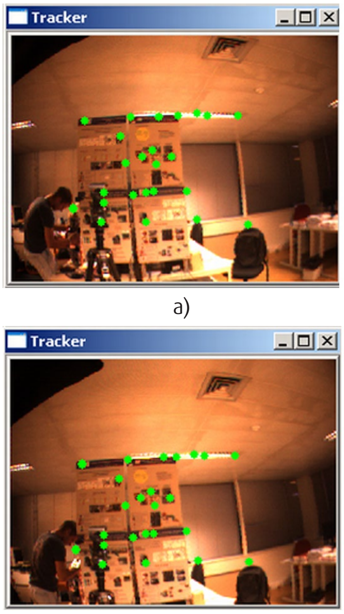

a)

b)

Figure 5 · Right eye image grabbed with tracked points, before (a) and after (b) the tilt motion.

The same procedure can be applied to the left eye, so that both eyes are calibrated. Once the eyes are calibrated, then we can calibrate the neck pan angle using the same method. This is illustrated in Fig. 6.

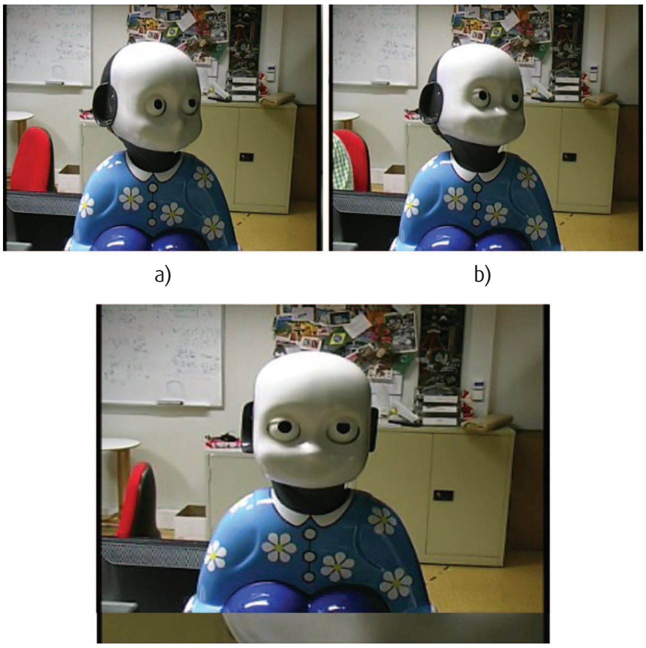

Figure 6 · Calibration procedure for the neck pan. The robot starts up with significant offset in the neck pan joint, with respect to its canonical position (a). A motion of 10 degrees is applied to the neck tilt (b). The proposed method estimates the initial pan angle and compensates it (c).

Notice that, in the beginning, the neck pan angle has a large offset to its canonical position. In the calibration procedure we have applied a tilt motion of 10 degrees up. Then the neck pan angle was calculated and the head set to its canonical position. In the end, the eyes are still a bit tilted, because the tilt joint was not calibrated in this test.

For a preliminary quantitative evaluation of the method's performance, we have initialized the system in a known position, and measured the pan angle using the proposed method, applying tilt motions of different amplitudes. Then, the pan was set to its estimated zero position, and the process iterated in a series of steps. The results are represented in Fig. 7, where it can be observed that the axis reaches a close vicinity of the correct position immediately after the first iteration. Then, in the remaining steps, the estimated positions are always kept within a 5 degrees range to the ideal zero position. In any case we had always good systematic results, even for small angular prospective motions.

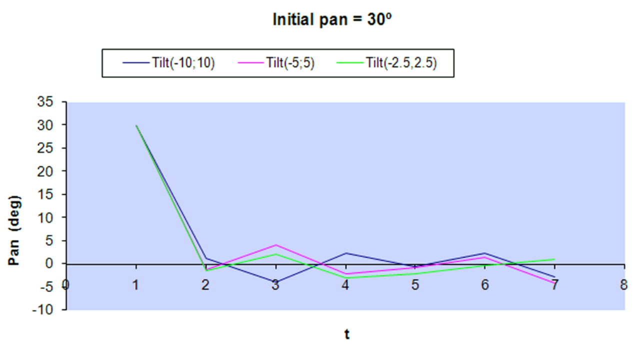

Figure 7 · Evolution of the pan angle for an iterative application of the calibration method. Different colors represent different prospective motion amplitudes.

The algorithm was tested many times for different circumstances. We came to the conclusion that results depend mostly on the quality of tracking. We had the best results while putting in front of robot a chessboard pattern, where point matches are very reliable. In other cases, there were a lot of points along edges in the image. This may lead to tracking drifts, and results may degrade a bit.

5. CONCLUSIONS

We have presented the principles for a vision based automatic calibration procedure for determining the initial unknown angles of pan-tilt kinematic structures. The method is based on the computation of the homography induced by the rotation of the tilt axis. A set of points is tracked in the images before and after the prospective motions. A robust estimation architecture allows the estimation of the homographies from the tracked points, even in the presence of tracker failures (outliers). By relating the homography entries with the unknown initial angles, it is possible to estimate them reliably from the visual measurements.

In future work we will further characterize the precision of the method as a function of the initial conditions and motion amplitudes. We also aim at investigating how errors propagate if the method is to be employed in longer kinematics chains. Finally we will study the combination of information from both cameras to improve precision in homing the neck joints

ACKNOWLEDGEMENTS

This work was supported by the European Commission, Project IST-004370 RobotCub, and by the Portuguese Government Fundação para a Ciência e Tecnologia (ISR/IST plurianual funding) through the PIDDAC Program funds, and through project BIO-LOOK, PTDC/EEA-ACR/71032/2006.

REFERENCES

[1] Beira, R., Lopes, M., Praça, M., Santos-Victor, J., Bernardino, A., Metta, G., Becchi, F. and Saltarén, R. Design of the Robot-Cub (iCub)Head, Proc. IEEE International Conference on Robotics and Automation, ICRA, Orlando, May, 2006.

[2] Metta, G., Fitzpatrick, P. and Natale, L. YARP: yet another robot platform, International Journal on Advanced Robotics Systems, Special Issue on Software Development and Integration in Robotics, March, 2006.

[3] Galassi, M. et al, GNU Scientific Library Reference Manual – Revised Second Edition, ISBN0954161734

[4] Intel Research. Open Computer Vision Library. http://www.intel.com/research/mrl/research/opencv/

[5] Hartley, R. and Zisserman, A. Multiple View Geometry in Computer Vision. Cambridge University Press 2000.

[6] Fischler, M. A. and Bolles, R. C. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. Comm. of the ACM, Vol 24, pp 381-395, 1981.

[7] Shi, J. and Tomasi, C. Good Features to Track. IEEE Conference on Computer Vision and Pattern Recognition, pages 593-600, 1994.

[8] Bouguet, J. Pyramidal Implementation of the Lucas Kanade Feature Tracker, Open CV Documents, Microprocessor Research Labs, Intel Corporation, 1999.

[9] Basri, R., Rivlin, E. and Shimshoni, I. Visual Homing: Surfing on the Epipoles. ICCV'98.

[10] Jin, Y. and Xie, M. Vision Guided Homing for Humanoid Service Robot. ICPR'00.