

Multi-modal tracking using texture changes

Christopher Kemp *, Tom Drummond

Department of Engineering, University of Cambridge, Cambridge CB1 2PZ, UK

Received 21 February 2005; received in revised form 15 November 2006; accepted 18 December 2006

Abstract

We present a method for efficiently generating a representation of a multi-modal posterior probability distribution. The technique combines ideas from RANSAC and particle filtering such that the 3D visual tracking problem can be partitioned into two levels, while maintaining multiple hypotheses throughout. A simple texture change-point detector finds multiple hypotheses for the position of image edgels. From these, multiple locations for each scene edge are generated. Finally, we determine the best pose of the whole structure. While the multi-modal representation is strongly related to particle filtering techniques, this approach is driven by data from the image. Hence the resulting system is able to perform robust visual tracking of all six degrees of freedom in real time. Real video sequences are used to compare the complete tracking system to previous systems.

© 2007 Published by Elsevier B.V.

Keywords: RANSAC; Texture; Real-time; Tracking; Multi-modal

1. Introduction

Robust real-time visual tracking has important applications such as the guidance of robots or autonomous vehicles. Our ultimate goal is the control of a miniature aerial vehicle (MAV), such as a model helicopter, and this provides particularly challenging problems. To control the flight it is necessary to have a real-time estimate of the vehicle's position and orientation (pose). In previous work with large model helicopters, e.g. [1], this has been achieved using a combination of GPS measurements, inertial sensors and visual information. Our aim is to control much smaller indoor vehicles which have greatly reduced payload capacities. Accurately locating the vehicle using only the information from a sub-miniature transmitting video camera would therefore be a great advantage.

Visual tracking systems aim to follow motion in a sequence of images from a video camera. If the video camera is moving through a static world, this can be used to recover the motion of the camera in real-time. Although

a prototype MAV system is not yet complete, initial experiments suggest that images obtainable are likely to contain large pixel noise and suffer from substantial motion blur. Further, MAVs are particularly unstable platforms and hence the tracking system must be able to handle large motion disturbances [Table 1](#).

To address these problems, we propose a novel multiple hypothesis approach which can reliably track polyhedral surroundings for which an accurate 3D edge model is known. This paper makes two contributions. In [Section 2](#) we show how ideas from RANSAC can be combined with particle filtering to generate multi-modal posterior distributions. A key benefit of this method is that it permits the tracking problem to be decomposed into two steps. Firstly, given noisy edgel locations, a multi-modal posterior for the location of each scene edge can be generated. Secondly, these can be used to generate a multi-modal posterior for the camera pose. The system requires multiple edgels to be found along image scan-lines and so in [Section 3](#), we describe the second contribution: a method of finding multiple texture change-points in a pixel sequence. The effectiveness of the resulting real-time tracking system is discussed in [Section 4](#).

* Corresponding author. Tel.: +44 785 526 2758.

E-mail address: ck231@cam.ac.uk (C. Kemp).

Table 1

Results from tests on three similar corridor sequences, showing the number of tracking failures for three different schemes

Seq.	No. of frames	No. of failures		
		Boujou	Earlier edge based	Multi-modal RANSAC
1	623	18	13	0
2	763	35	21	6
3	692	22	8	3

1.1. Previous work

Existing real-time tracking systems can be split into two groups: those which rely on a predetermined model of the tracked scene and those which learn the surroundings as the sequence progresses. Some very encouraging progress has been made recently in the latter group [2] and since the construction of models is costly, systems which automatically create or adapt models are advantageous. However, in areas where the scene is reasonably static, the use of prior knowledge (in the form of a model) can only be beneficial. In particular, model based systems can more easily reject unmodelled features and so are more tolerant of non-static elements in the images, such as passing humans.

Many model based tracking schemes work by matching point features. Some methods sequentially match features from adjacent frames e.g. [3]. Recent contributions [4] have largely eliminated both drift and jitter through the use of inter-frame matching combined with precomputed keyframes. Whilst many scenes can be successfully tracked using point features, some indoor areas where our MAV may be used are surprisingly free of reliable point features. Also, point detectors can have difficulties with images with repetitive textures such as carpets or tiles and struggle with noisy and blurred images.

Another technique is to match image edges. Edge detection is more robust to illumination and aspect changes and it is more straightforward to detect edges in the presence of noise and blur. Typically however it is difficult to match image edges based on appearance and so a common assumption [5–8] is that closer correspondences (to a predicted position) are more likely to be correct. However,

if there is a high density of similar features, or large unpredictable image motions occur, this assumption is violated. Although robust estimators are used to tolerate a small number of false correspondences, in these conditions tracking often converges on an incorrect local minimum. It is also possible [9] to work directly with image intensities rather than first performing edge extraction. However, the same problems with local minima apply.

A number of techniques exist to deal with local minima. Particle filtering has proved successful for many applications [10,11]. However, when tracking motion in all six degrees of freedom, it is difficult to achieve the required robustness in real-time. Alternative schemes which also consider multiple hypotheses have been used. In [12], the multi-modal representation used is a piecewise combination of Gaussian distributions and is used to track rapidly moving human figures, while in [13] image corners are used with a discrete, tree-like, representation of the multiple hypotheses. RANSAC [14] based approaches have also been used [5] and numerous extensions to ransac have been proposed, in particular the idea of using RANSAC to initialise a set of particles is employed by [15] as part of a structure from motion system.

1.2. Overview of the approach

The proposed system follows the same basic structure as many frame-to-frame tracking systems, such as the early method described in [6]. For each new frame a 3D model of the surroundings is rendered into the imaging plane at the pose obtained from the previous frame. The image is next searched to identify the new location of the rendered model edges. This is performed using 1D searches along regularly spaced scan-lines, perpendicular to each rendered model edge. Finally, an optimisation is performed, adjusting the six pose parameters to align the rendered edges with those found in the image. These stages are shown in Fig. 1.

The system proposed in [5] used RANSAC to improve edge-based tracking by robustly estimating the location of each model edge in the image. Here, we extend this by admitting multiple hypotheses for each line, as described in the next section. This gives the resilience to local minima demonstrated by particle filtering approaches whilst still providing real-time operation.




Fig. 1. An overview of the system tracking a carpet tile and screen, despite occlusion. From the left: for each new frame the model is rendered in the predicted location. Next, perpendicular searches locate matches for each model edge in the image. Finally, the best overall pose is found.

2. Multi-modal posterior representations

The Random Sample Consensus (RANSAC) paradigm, introduced by Fischler and Bolles [14], has been widely used throughout the computer vision field to obtain robust estimates of model parameters. The prototypical example of its use is to obtain the parameters of a straight line which best fits a set of noisy measurements, M , containing outliers. The parameters describing the straight line through two measurement points selected at random are obtained. A consensus count for this hypothesis is then found by comparing, to a threshold value, the distance from the line to all the measurement points. The process is repeated N times and the hypothesis with the highest consensus is finally selected.

The consensus system proposed in [14] can be expressed as:

$$C(\theta) = \sum_{m \in M} f(\text{dist}(m, \theta)), \quad (1)$$

where $\text{dist}(m, \theta)$ is a function giving the distance from a measurement point m to the straight line defined by θ and:

$$f(d) = \begin{cases} 0 & (d < t) \\ 1 & (d > t) \end{cases}, \quad (2)$$

where t is the consensus range or threshold. Defined in this way, the hypothesis with the smallest consensus score is chosen. In [16] it was observed that the performance of the system is improved if, at no extra computational cost, this last function is replaced with:

$$\tilde{f}(d) = \begin{cases} d & (d < t) \\ t & (d > t) \end{cases}. \quad (3)$$

When the consensus score is expressed in this way, RANSAC is being used to find the minimum of a function defined continuously over model parameters. [16] shows how (under certain assumptions) $\tilde{f}(d)$ can be used to obtain a consensus score which is the negative log likelihood that the hypothesis is correct. From this the posterior probability $P(\theta)$ can be easily calculated:

$$P(\theta) = K e^{-C(\theta)/\sigma}. \quad (4)$$

The approach proposed here borrows an idea from condensation and retains *all* of the hypotheses tested. Each is then treated as a particle in the representation of the posterior distribution. Suppose that a set of measurements is obtained which contains noisy points from more than one straight line, along with some outliers. RANSAC applied to this data set would find the parameters of only the line with the most corresponding measurement points (given a sufficiently large N). Instead, if the consensus scores from all N RANSAC hypotheses are retained, these can be used to give an approximation to the multi-modal posterior. This representation is beneficial if the next task is to draw samples from the posterior since this can be approximated by drawing from the stored hypotheses. A good approximation can be obtained with a relatively

low N since there are typically many proposal points at or near the modes of P . However, this means that it is *not* correct to simply select hypothesis n with probability $p_n = P(\theta_n)$.

If a posterior distribution to be represented is $P(\theta)$, and given a set of points $\Theta = \theta_1 \dots \theta_N$ drawn from a (different) proposal distribution $Q(\theta)$, then draws from the posterior distribution can be approximated by resampling from Θ , where the probability of selecting sample n is:

$$P_n \propto \frac{P(\theta_n)}{Q(\theta_n)}. \quad (5)$$

Provided that $Q(\theta) > 0$ for all θ where $P(\theta) > 0$ this process converges to sampling from P as $N \rightarrow \infty$. In the case of multi-modal RANSAC, the proposals are those generated by random sampling of data tuples created by a process driven by P . The probability of obtaining the sample point, $Q(\theta_n)$, is a function of the consensus set (the number of data points within consensus range of θ_n). If the size of the consensus set is c_n and the number of data samples needed to form a hypothesis is j , then Q is given by:

$$Q(\theta_n) = k^{c_n} C_j = k \frac{c_n!}{j!(c_n - j)!}, \quad (6)$$

provided that the same consensus set is obtained from any j -tuple subset. Although this is not strictly the case, in practice the approximation works well. Hence when sampling from a multi-modal representation of a posterior generated by RANSAC, each hypothesised sample, θ_n should be chosen with a probability, p_n , given by:

$$P_n = K \frac{P(\theta_n)}{c_n C_j}, \text{ with } K \text{ chosen such that } \sum P_n = 1. \quad (7)$$

2.1. Application to finding straight edges in an image

Given an initial estimate for the location of an edge in an image, the task is to generate the posterior distribution of the true position of that edge, given the image data. Since the true edge is expected to be close to the initial prediction, it is convenient to express this posterior as a distribution over motion from the predicted location. Due to the aperture problem, motions tangential to the edge direction are very difficult to detect accurately and so only motions normal to the edge direction are modelled. The posterior distribution is therefore defined over two model parameters r_1 and r_2 , which describe the motion of the two endpoints of the line in a direction normal to the line (see Fig. 2). First, edgels are obtained, either at points of maximal intensity gradient, or using the scheme described in Section 3. The ideas proposed in the previous section are then applied to find straight lines through these edgels.

The required posterior distribution describes the probability density that the corresponding straight edge in the image is at $\theta = \{r_1, r_2\}$. An image edge is modelled as causing edgels which are displaced from a straight line by additive Gaussian noise. If $d_{n,s}$ is the distance along each




Fig. 2. (a) A closeup of Fig. 1 showing a rendered edge (black line), the two model parameters (r_1 and r_2), the detected edgels (black dots), and samples from the multi-modal representation (white lines), (b) contours of the true multi-modal posterior computed using a dense search, (c) the sampled representation at the same scale.

perpendicular scanline s , from the intersection with the hypothesis edge θ_n to the nearest edgel, a close approximation¹ to the negative log likelihood is:

$$C(\theta_n) = \text{const} + \sum_s \tilde{f}(d_{n,s}^2), \quad (8)$$

with $\tilde{f}()$ as in (3). Hence the posterior is given by:

$$P(\theta_n) = K e^{-\frac{C(\theta_n)}{2\sigma^2}}. \quad (9)$$

In this case the N samples are obtained by randomly selecting two edgels and calculating the parameters that describe the straight line through these two points. The multi-modal representation is completed by calculating and storing $\Theta_n = \{\theta_n, P(\theta_n)\}$ for each. Since in this case $j=2$, samples can then be drawn from the representation by selecting θ_n with a probability proportional to:

$$p_n = \frac{P(\theta_n)}{c_n(c_n - 1)}, \quad (10)$$

where c_n is the corresponding consensus count as defined for (6).

The resulting system is demonstrated in Fig. 2 which compares the full posterior distribution to that generated by sampling the RANSAC representation, with $N = 100$.

2.2. Determining camera pose

Determination of camera pose relative to a rigid structure is a 6 DoF problem. Given a 3D edge model of the structure in question, and a set of candidate edgel matches for sample points on this model, then RANSAC could be applied to these, selecting six samples to generate each hypothesis. However, this is very inefficient in the presence of significant noise (for example when there are many candidate matches for each edgel as in our case). One of the key benefits of the approach presented here is that it

permits a hierarchical partitioning of this problem and thus, we find pose in two stages:

- (1) The 3D edge model is rendered and for each visible edge segment and candidate edgel matches are found. The techniques of Section 2.1 are then applied to find a multi-modal posterior for matches to each visible edge segment.
- (2) The candidate matches for each edge segment are then used in a second stage RANSAC which selects matches for three edge segments to generate each pose hypothesis.

The main advantage of this two stage approach is that the very large number of outliers is handled in the first stage where only two samples are needed to form each hypothesis (rather than six if pose were to be solved for directly). This then only emits candidates for which significant support exists and the second stage is able to use these to rapidly determine pose.

Thus, having found a set of hypothesised correspondences for each model edge, the second level RANSAC is used to find a new camera pose which maximises consensus of the original edgels. Three model edges are selected at random and for each a corresponding image line is chosen according to (10). This is usually sufficient to constrain a new hypothesised pose. It is straightforward to obtain a Jacobian matrix which describes the rate of change of each line parameter, r_i , with respect to each pose parameter, α_j , about the current pose $\underline{\alpha}$. Hence a linear approximation to the new hypothesised pose is given by:

$$\hat{\underline{\alpha}} = \underline{\alpha} + J^{-1} \underline{r} \quad \text{where } J_{ij} = \frac{\partial r_i}{\partial \alpha_j} \quad (i = 1 \dots 6, j = 1 \dots 6). \quad (11)$$

The consensus score for the second stage RANSAC is found using a full Jacobian matrix to give a linear prediction for the new parameters of all the model lines. The probability density of each new line is evaluated from the original edgels using (9) and the product of these then

¹ This approach fixes the number of inliers to be constant and hence the inlier/outlier ratio has only a small effect on the true negative log likelihood.

gives the overall pose posterior. The second stage RAN-SAC is performed L times and the pose with the highest posterior is retained as the estimate of the camera pose for that frame.

2.3. Relationship with particle filtering

As mentioned previously, the multi-modal representation proposed here is strongly related to that employed in particle filtering systems such as [11]. In both cases the true continuous posterior distribution is represented by measurements at a set of discrete sample points. In CONDENSATION, the assumption is that it is computationally cheap to measure the posterior at these sample points. Hence a relatively dense set of sample points can be generated at locations of high prior probability.

Unfortunately with a relatively complicated edge model, it is only possible to measure the posterior around 500 times per frame. Even if the prior for the frame has only one single peak, positioning 500 particles around this peak in six degrees of freedom gives just two or three measurements in each direction. This severely limits the motion disturbances which can be tolerated by traditional particle filtering schemes.

The system proposed here instead uses sample points which are directly suggested by the image data. Hence the particles are concentrated into areas which are likely to be near the peaks of the posterior. This significantly reduces the number of particles needed and allows larger motion disturbances to be tracked successfully.

3. Finding texture change-points in a 1D line search

The scenes that our system is exposed to frequently contain a significant amount of visual texture and early experiments showed that at a conventional edgel detector which fires on significant local maxima of intensity gradient (e.g. [6,11]) is inadequate for our use. In recent work [17] a detector was described which finds the single most probable location of a texture change. This texture operator has the important advantage of being extremely fast, but it can only yield a single candidate edgel along each scanline and thus (with non-trivial frequency) returns a false match.

Here this is extended to allow multiple possible change-points. This allows the tracking system to operate in textured environments yet adds little extra computational cost.

In [17] a tracking system similar to [6] was proposed which used a 1D texture change-point search instead of a simple gradient based edgel detector. In this paper, pixel intensities are first grouped into regularly spaced intensity bins. A known zeroth-order texture generating process T_1 is then defined as one where pixel intensities are independently drawn from a probability distribution over I intensity bins ($T_1 = \{p_i; i = 1 \dots I\}$) (i.e. independently from a histogram). Hence given a sequence of N binned pixel

intensities $S_0^N = (s_0, s_1, \dots, s_{N-1})$, the probability of obtaining that sequence, given process T_1 is:

$$P(S_0^N | T_1) = \prod_{n=0}^{N-1} p_{s_n}, \quad (12)$$

The key result from [17] is that if the texture process is instead unknown, but all possible texture processes are equally likely (i.e. there is a uniform prior on T), then given a sequence of n pixel intensities $S_0^n = (s_0, s_1, \dots, s_{n-1})$ from T , the expected probability of a further sample, s_n from T is given by:

$$E(p_{s_n} | S_0^n) = \frac{O_{s_n} + 1}{n + I}, \quad (13)$$

where there are o_j occurrences of the symbol j in the sequence S_0^n . Hence the probability of the entire sequence given a single texture is:

$$P(S_0^N) = \prod_{i=0}^{N-1} E(p_{s_i} | S_0^i) = \frac{(I-1)! \prod_{i=1}^I (o_i + 1)!}{I + N - 1}, \quad (14)$$

In [17] this result was used to find the most probable location for a single texture change-point along a scanline and an efficient algorithm for this was presented. In this work however, multiple change-points must be considered. Instead of maximising the probability of the location of the single change-point given a sequence of intensities, we must find the set of change-points with maximum probability. If the sequence of N pixels, S , is modelled by $m-1$ distinct textures between m change-points, $M = (c_1, c_2, \dots, c_m)$:

$$P(S|M) = \prod_{j=1}^m P(S_{c_j}^{c_{j+1}}). \quad (15)$$

Finally using Bayes' rule we maximise:

$$P(M|S) = \frac{P(S|M)P(M)}{K}. \quad (16)$$

For this work $P(M)$ was chosen simply as $P(M) = \lambda^m$, with λ constant and < 1 .

Evaluating (16) for all 2^N possible combinations of change-points would be computationally prohibitive. Fortunately it is possible to use a dynamic programming algorithm to improve efficiency greatly. The most probable set of change-points can be determined by considering the location of the last change-point in the sequence, c_m . Suppose that $\text{argmax}_k P(S_1^k | M')$ is known for all $k < l$. The most probable position of c_m is then:

$$\begin{aligned} c_m &= \text{argmax}_c P(S_0^l | \text{last change-point at } c) \\ &= \text{argmax}_c \left[\max_{M'} P(S_0^c | M') P(S_c^l | \emptyset) \right], \end{aligned} \quad (17)$$

since the probability of the sequence after c_m is independent of the sequence up to c_m . \emptyset is the empty set implying no change-points. Hence:

$$\begin{aligned} \operatorname{argmax}_M P(S_0^l|M) &= \{l\} \cup \operatorname{argmax}_{M'} P(S_0^{c_m}|M') \\ \text{and } \max_M P(S_0^l|M) &= \max_{M'} P(S_0^{c_m}|M') P(S_{c_m}^l|\emptyset). \end{aligned} \quad (18)$$

By induction these expressions can be evaluated sequentially for $l = 0 \dots N$.

This means that, for $l = 0$, $P(S_0^0|0, 1)$ is evaluated using s_0 . At the second step both $P(S_0^1|0, 2)$ and $P(S_0^1|0, 1, 2)$ must be evaluated using the value from the first step and s_1 . Using these values, there are four potential sets of change-points to evaluate at the third step. However, the change from $P(S_0^1|0, 2)$ to $P(S_0^2|0, 2, 3)$ is the same as that from $P(S_0^1|0, 1, 2)$ to $P(S_0^2|0, 1, 2, 3)$ since by (15), the probability of the sequence after a change-point at 2 is independent of the sequence before it. Hence the most probable of $P(S_0^1|0, 2)$ and $P(S_0^1|0, 1, 2)$ will lead to the most probable of $P(S_0^2|0, 2, 3)$ and $P(S_0^2|0, 1, 2, 3)$. So at the third step we can simply pick the highest probability from the second step, and only evaluate three new probabilities. This means that when processing the n th pixel we must only evaluate n possibilities, giving a total of $N^2/2$ operations for the entire sequence. The complete operation is shown in Algorithm 1.

Algorithm 1 Finding texture change-points along a 1D scanline

```

dim Observations[NUM_PIXELS] [NUM_BINS]
dim LogLikelihoods[NUM_PIXELS],   Changepoints
[NUM_PIXELS]
BestLogLikelihood = 0
for i = 1 .. NUM_PIXELS do // Loop for each pixel in
                             the sequence
  for j=1 .. NUM_BINS do // Set up a new node with a
                           cleared histogram
    Observations[i][j] = 1 // Seed with 1 sample per bin
  end for
  //and start with the best log likelihood to that change-
  //point plus prior 'penalty'
  LogLikelihoods[i] = BestLogLikelihood - log( $\lambda$ )
  BestLogLikelihood = BIG_NUM
  for k = 1 .. i do // Loop round the nodes
    //update the node's log likelihood given the new
    //pixel using [13]
    LogLikelihoods[k] += log( $\sum$  Observations[k][j]) -
    log(Observations[k][Pixels[i]])
    Observations[k][Pixels[i]] += 1
    //and store the best ready for the next outer loop
    if LogLikelihoods[k] < BestLogLikelihood then
      BestLogLikelihood = LogLikelihoods[k]
      Changepoints[i] = k
    end if
  end for
end for
//The optimal changepoints can now be found by working
back:
//c1=Changepoints[NUM_PIXELS],   c2=Changepoints
[c1], etc.

```

Whilst this is not a major computational load, a second very simple speedup is possible. Since 16 will never be maximal if there is a change-point between pixels of the same binned intensity, there is no need to consider such cases. This offers a massive speed improvement in areas of the image which have a relatively constant intensity.

4. Results

Fig. 1 showed a simple example of successful tracking which demonstrates several benefits. Firstly, the system is able to track in the presence of simple textures, such as carpets. Secondly, the multi-modal posterior representation means that the system is able to consider multiple correspondences. In Fig. 1, the chair produces a stronger edge than the carpet tile but, since both edges appear in the multi-modal posterior, the system sticks to the correct edge since this provides a better overall match. Finally, the scheme is relatively tolerant to occlusion. Although both the chair and the wastebasket in Fig. 1 are unmodelled clutter, the system (given a large enough L) should find the correct pose as three of the modelled edges are visible in the image. Of course, clutter will reduce the overall pose posterior score, but typically the best score will be the correct pose.

The system was also tested on two full video sequences, one of a staircase and one of a corridor. Although both sequences were prerecorded, processing was performed at 10 and 30 fps, respectively (on a single processor 2.4 GHz Pentium). In both cases the predicted location for each frame was simply that found in the previous frame.

4.1. Staircase sequence

The staircase sequence (Fig. 3) demonstrates the system's ability to cope with textured surfaces as well as rapid motion (leading to large inter-frame changes) and significant orientation changes. Attempts to track this sequence using our earlier edge-based system [8] fail almost instantly as shown in Fig. 4. Tracking also fails if a simple gradient based edge detector is used instead of the scheme proposed in Section 3. Since the speckled steps produce a huge number of hypothesis edgels, a very large number of RANSAC hypotheses (large N and L) would be needed to track successfully. Using the scheme proposed in Section 3 makes the task feasible and the sequence tracks correctly.

4.2. Corridor sequence

The corridor sequence uses relatively poor video, such as might be expected from a miniature transmitting camera. Some clips showing the sequence being successfully tracked are shown in Fig. 5. Tolerance of significant pixel noise and blurring, and again robustness to occlusion, is demonstrated.

The corridor sequence was also tested on two other tracking systems, a demonstration version of 'boujou'




Fig. 3. Frames 1, 50, 100, 150 from a staircase video, being correctly tracked by the multi-modal RANSAC system.




Fig. 4. Left, a closeup of frame 10 from the staircase sequence showing the single hypothesis system failing to find the correct edges. Right, using multi-hypothesis RANSAC also fails if a simple gradient based edge detector is used.



Fig. 5. Frames 1, 200, 400, 600 from a corridor sequence, being correctly tracked by the multi-modal RANSAC system.




Fig. 6. Left, a closeup showing typical blur levels. Right, reliable inter-frame point matching is difficult with this level of blur.

and our earlier edge based system [8]². ‘boujou’ is a commercial product (www.2d3.com) which performs

² While the original version performed a single iteration per frame on a 200 MHz machine, we now aid convergence by performing many tracking iterations per frame.

offline bundle adjustment of tracked interest points. For a simple quantitative comparison we have chosen to measure the number of tracking failures suffered by each method on three different sequences from the same corridor. After each failure the pose was manually reset.



Fig. 7. Two examples of our earlier single hypothesis edge tracking system failing to track correctly. These frames are tracked correctly with the multi-modal RANSAC approach.

The blur and pixel noise present in the images means that reliable point features are difficult to obtain and match, as shown in Fig. 6. Hence boujou was generally unable to maintain track for more than a few frames at a time.

Fig. 7 shows two typical occasions where our earlier, single hypothesis method was unable to track correctly. In both cases the tracker has converged on an incorrect local minimum since modelled edges have locked onto closer parallel edges in the image. Only by considering multiple hypotheses for each edge can these failures be avoided.

Failures of the multi-modal system occurred for three reasons. Firstly, the second sequence contains 18 frames which are largely obscured by a passer-by. Since the camera was rotating at this time, no method tested was able to maintain track during this period of occlusion. Most importantly, some views of the corridor scene are constrained rather poorly using just three correspondences between image lines and model edges. Hence, even with exhaustive RANSAC samples, the correct pose will never be obtained. This failure mode is the topic of future work. Finally, when the numbers of RANSAC hypotheses (N and L) are chosen to allow real time operation, there remains a limit on the magnitude of motion disturbance which can be tolerated and failures still occur if the camera is moved or accelerated too fast.

5. Conclusions

This paper has presented a method for generating a representation of a multi-modal distribution. This is particularly beneficial when partitioning a problem such that the multiple hypotheses from one stage provide samples for further processing. The method has been successfully applied to the task of locating straight edges in an image using the results from a novel texture change-point detector and subsequently to perform visual tracking of known polyhedral objects. The system shows a significant improvement in motion robustness over

previous techniques. Additionally, due to the data driven approach, it can be used to track all six degrees of freedom in real time.

References

- [1] J. Dittich, Design and integration of an unmanned aerial vehicle navigation system, Master's thesis, Georgia Institute of Technology, Atlanta, GA 30332 (July 2002).
- [2] A.J. Davison, Real-time simultaneous localisation and mapping with a single camera, in: Proceedings of the Ninth International Conference on Computer Vision, 2003, pp. 1403–1401.
- [3] G. Simon, A. Fitzgibbon, A. Zisserman, Markerless tracking using planar structures in the scene, in: Proceedings of the International Symposium on Augmented Reality, 2000, pp. 120–128.
- [4] L. Vacchetti, V. Lepetit, P. Fua, Fusing online and offline information for stable 3D tracking in real-time, in: Proceedings of the International Conference on Computer Vision and Pattern Recognition, vol. 2, 2003, pp. 241–248.
- [5] M. Armstrong, A. Zisserman, Robust object tracking, in: Proceedings of the Asian Conference on Computer Vision, vol. I, 1995, pp. 58–61.
- [6] C. Harris, C. Stennett, RAPID – a video rate object tracker, in: Proceedings of the British Machine Vision Conference, 1990, pp. 73–77.
- [7] D. Lowe, Robust model-based motion tracking through the integration of search and estimation, *International Journal of Computer Vision* 8 (2) (1992) 113–122.
- [8] T. Drummond, R. Cipolla, Real-time visual tracking of complex structures, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (7) (2002) 932–946.
- [9] E. Marchand, P. Bouthemy, F. Chaumette, V. Moreau, Robust real-time visual tracking using a 2D–3D model-based approach, in: Proceedings of the International Conference on Computer Vision, vol. 1, 1999, pp. 262–268.
- [10] J. Deutscher, A. Blake, I. Reid, Articulated body motion capture by annealed particle filtering, in: Proceedings of the International Conference on Computer Vision and Pattern Recognition, vol. 2, 2000, pp. 126–133.
- [11] M. Isard, A. Blake, CONDENSATION – conditional density propagation for visual tracking, *International Journal of Computer Vision* 29 (1) (1998) 5–28.
- [12] T. Cham, J. Rehg, A multiple hypothesis approach to figure tracking, *Computer Vision and Pattern Recognition* (1999) 239–245.
- [13] I. Cox, S. Hingorani, An efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18 (2) (1996) 138–150.

- [14] M. Fischler, R. Bolles, Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, *Communications of the A.C.M.* 24 (6) (1981) 381–395.
- [15] P. Torr, C. Davidson, IMPSAC: synthesis of importance sampling and random sample consensus, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (3) (2003) 354–364.
- [16] P. Torr, A. Zisserman, MLESAC: a new robust estimator with application to estimating image geometry, *Computer Vision and Image Understanding* 78 (1) (2000) 138–156.
- [17] A. Shahrokni, T. Drummond, P. Fua, Texture boundary detection for real-time tracking., in: *Proceedings of the European Conference on Computer Vision*, vol. 2, 2004, pp. 566–577.