



Tracking of 2D or 3D Irregular Movement by a Family of Unscented Kalman Filters

Junli Tao* and Reinhard Klette, *Member, KIICE*

Department of Computer Science, Auckland University, Auckland 1010, New Zealand

Abstract

This paper reports on the design of an object tracker that utilizes a family of unscented Kalman filters, one for each tracked object. This is a more efficient design than having one unscented Kalman filter for the family of all moving objects. The performance of the designed and implemented filter is demonstrated by using simulated movements, and also for object movements in 2D and 3D space.

Index Terms: Unscented Kalman filter, Detection-by-tracking, 2D/3D tracking

I. INTRODUCTION

Object tracking is an important task for many applications, such as for robot navigation, surveillance, automotive safety, and video content indexing. Based on trajectories obtained through tracking, some advanced behaviour analysis can be applied. For instance, a pedestrian's trajectory can be analysed to warn a driver if the trajectories of the vehicle and of the pedestrian are potentially intersecting.

For multiple object tracking, tracking-by-detection methods are the most popular algorithms. A detector is used in each image frame to obtain candidate objects. Then, with a data-association procedure, all the candidates are matched to the existing trajectories as known up to the previous frame. Any unmatched candidate starts a new trajectory. Since there is no perfect detector that detects all objects without any false positives and false negatives, sometimes objects are missed (i.e., they appear in the image but are not detected), or background windows are incorrectly detected as being objects. Such false-positive or false-negative detections increase the difficulty of tracking.

Occlusion by other objects or the background is one of the main reasons for detection to fail and it also increases the difficulty of tracking (e.g., identity switch). Some algorithms [1, 2] propose tracking objects in the 2D image plane. The occlusion problem is handled either using part detectors and tracking detected body parts, or adopting instance-specific classifiers to improve performance of data assignment. However, tracking in the 2D image plane increases the ambiguity of data association. A tall person nearby, and a small person far away, for example, may appear very close to each other in the image, and, possibly, in some frames the tall person occludes the small person. But they are actually several meters away from each other. Thus, often, and also in this paper, stereo information is adopted to improve the tracking performance [3-5], and multiple pedestrians are tracked in 3D coordinates.

Tracking objects with irregular movements in 3D space is a challenging task due to the totally unknown speed and direction. In this paper, the application of an unscented Kalman filter (UKF), which can also handle nonlinear—in fact, fully irregular trajectories in 3D space, is demonstrated. For the

Received 20 June 2012, Revised 10 July 2012, Accepted 18 July 2012

*Corresponding Author E-mail: jtao076@aucklanduni.ac.nz

Open Access <http://dx.doi.org/10.6109/jicce.2012.10.3.307>

print ISSN:2234-8255 online ISSN:2234-8883

© This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/3.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © The Korea Institute of Information and Communication Engineering

original paper on UKF see [6]. Similar work is proposed in [5]. However, instead of modelling the motion of the vehicle and the pedestrians separately, we straightforwardly model the relative motion between them, and no ground plane is assumed, so that objects moving with 6 degrees of freedom can be tracked properly. Different types of models are tested and compared in both simulation and real sequences.

II. RELATED WORK

Multiple object tracking has attracts a great deal of attention recently in computer vision research. Today, an update of the review [7] from 2006 should also include work such as in [2-4, 8-13].

Kalman filters (KF) have been extensively adopted to deal with tracking tasks. A KF is a recursive Bayesian filter, firstly, using motion information to predict the possible position, followed by fusing the observation (detection) and predicted position. A linear KF is used for tracking (e.g., [7]) when movement is such that linear models may be used for approximation. Obviously, a linear model is not suitable for most cases. The extended Kalman filter (EKF) was designed [14] for handling a nonlinear model by linearizing functions using the Taylor expansion extensively. For example, an EKF has been used for simultaneous localization and mapping [15], and for pedestrian tracking [16]. A particle filter was used to handle the task in [17]. Performance similar to an EKF is reported in [3].

The UKF can handle a nonlinear model by using the unscented transform to estimate the first and second order moments of sigma points, which represent the distribution of a predicted state and predicted observations, and it appears that the UKF does this better than the EKF [18]. Thus, in this paper, a UKF is used for tracking multiple, irregularly moving objects in 3D space, which is a highly nonlinear problem.

III. UNSCENTED KALMAN FILTER

The unscented transform (UT) is the core component that enables the UKF able to be able to handle nonlinear models. Let L be the dimensionality of the system state $x_{t-1|t-1}$ at time $t-1$. If the system noise (process noise \mathbf{Q} and measurement noises \mathbf{R}) is not additive noise, the state is augmented before UT. In our case, random acceleration is introduced as process noise; thus, the state augmented with a process noise vector, is denoted by

$$X_{t-1|t-1}^{(a)}$$

and called *vector* for short. The dimension of the augmented vector depends on the process model, which is illustrated in

Section IV. Let $x_{t|t-1}$ denote the predicted state at time t when passing $x_{t-1|t-1}$ through process function \mathbf{f} . Let $y_{t|t-1}$ be the predicted observation at time t when passing $x_{t|t-1}$ through observation function \mathbf{h} .

The UT works by sampling $2L+1$ sigma vectors $X_i^{(a)}$ in the augmented state space (following [6]), forming a matrix $\mathbf{X}^{(a)}$. The covariance matrix in augmented state space is denoted by $\mathbf{P}^{(a)}$. Let

$$\mathbf{P}_{t-1|t-1}^{(xx)}$$

be the state covariance matrix (i.e., describing dependencies between components of a state x). Formally,

$$\begin{aligned} X_0^{(a)} &= x_{t-1|t-1}^{(a)} \\ X_i^{(a)} &= x_{t-1|t-1}^{(a)} + \left(\sqrt{(L+\lambda)\mathbf{P}_{t-1|t-1}^{(a)}} \right)_i \\ &\quad \text{for } i = 1, 2, \dots, L \\ X_i^{(a)} &= x_{t-1|t-1}^{(a)} - \left(\sqrt{(L+\lambda)\mathbf{P}_{t-1|t-1}^{(a)}} \right)_{i-L} \\ &\quad \text{for } i = L+1, L+2, \dots, 2L \\ \mathbf{X}^{(a)} &= \begin{bmatrix} \mathbf{X}^{(s)} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{X}^{(n)} \end{bmatrix} \end{aligned}$$

where λ is a positive real, used as a scaling parameter. These sigma vectors can be passed through a nonlinear function (e.g., \mathbf{f} or \mathbf{h}) one by one, thus defining transformed (i.e., new) sigma vectors such as

$$\mathbf{X}_{t|t-1}^{(s)} \text{ or } \mathbf{y}_{t|t-1}$$

Means $x_{t|t-1}$ or $y_{t|t-1}$ and covariance matrices

$$\mathbf{P}_{t|t-1}^{(xx)} \text{ or } \mathbf{P}_{t|t-1}^{(yy)}$$

are obtained as follows; take \mathbf{h} for example:

$$\begin{aligned} y_{t|t-1} &= \mathbf{h}(\mathbf{X}_{t-1|t-1}^{(s)}) \\ y_{t|t-1} &\approx \sum_{i=0}^L W_i^{(m)} Y_i \\ \mathbf{P}_{t|t-1}^{(yy)} &\approx \sum_{i=0}^L W_i^{(e)} (Y_i - \mathbf{y}_{t|t-1})(Y_i - \mathbf{y}_{t|t-1})^T \end{aligned}$$

with constant weights $W_i^{(e)}$. Details are given in [6].

The UKF is illustrated as follows. At first we initialize the state $\mathbf{x}=\mathbf{x}_0$ and state covariances $\mathbf{P}^{(xx)} = \mathbf{P}_0^{(xx)}$. For the augmented vectors, let

$$\begin{aligned} \mathbf{x}^{(a)} &= (\mathbf{x}^T \mathbf{0}^T)^T \\ \mathbf{P}^{(a)} &= \text{diag}(\mathbf{P}^{(xx)}, \mathbf{Q}) \end{aligned}$$

where \mathbf{Q} denotes the process-noise covariance matrix. Details about \mathbf{Q} are given in Section IV. For $t \in (1, \dots, \infty)$, we calculate sigma vectors as follows:

$$\begin{aligned} \mathbf{x}_{t-1|t-1}^{(a)} &= (\mathbf{x}_{t-1|t-1}^{(a)}, \mathbf{x}_{t-1|t-1}^{(a)} + \gamma \sqrt{\mathbf{P}_{t-1|t-1}^{(a)}} \\ &\quad \mathbf{x}_{t-1|t-1}^{(a)} - \gamma \sqrt{\mathbf{P}_{t-1|t-1}^{(a)}}) \end{aligned}$$

where $\gamma = \sqrt{\lambda + L}$. The process update is as follows:

$$\begin{aligned} \chi_{t|t-1}^{(s)} &= f(\chi_{t-1|t-1}^{(s)}, \chi_{t-1|t-1}^{(n)}) \\ \mathbf{x}_{t|t-1} &= \sum_{i=0}^{2L} W_i^{(m)} X_i^{(s)} \\ \mathbf{P}_{t|t-1}^{(xx)} &= \sum_{i=0}^{2L} W_i^{(c)} (X_i^{(s)} - \mathbf{x}_{t|t-1})(X_i^{(s)} - \mathbf{x}_{t|t-1})^T \end{aligned}$$

We update the sigma vectors using

$$\begin{aligned} \mathbf{x}_{t|t-1}^{(s)} &= (\mathbf{x}_{t|t-1}, \mathbf{x}_{t|t-1} + \gamma \sqrt{\mathbf{P}_{t|t-1}^{(xx)}} \\ &\quad \mathbf{x}_{t|t-1} - \gamma \sqrt{\mathbf{P}_{t|t-1}^{(xx)}}) \\ y_{t|t-1} &= h(\mathbf{x}_{t|t-1}^{(s)}) \\ y_{t|t-1} &= \sum_{i=0}^{2L} W_i^{(m)} Y_i \end{aligned}$$

and update the measurement covariance matrix as follows:

$$\mathbf{P}_{t|t-1}^{(yy)} = \sum_{i=0}^{2L} W_i^{(c)} (Y_i - y_{t|t-1})(Y_i - y_{t|t-1})^T + \mathbf{R}$$

where \mathbf{R} is the assumed measurement noise covariance, depending on the observation model selected. Details are given in Section IV.

Altogether, the UKF is defined by

$$\begin{aligned} \mathbf{P}_{t|t-1}^{(xy)} &= \sum_{i=0}^{2L} W_i^{(c)} (X_i - \mathbf{x}_{t|t-1})(Y_i - y_{t|t-1})^T \\ K_t &= \mathbf{P}_{t|t-1}^{(xy)} (\mathbf{P}_{t|t-1}^{(yy)})^{-1} \\ \mathbf{x}_{t|t} &= \mathbf{x}_{t|t-1} + K_t (y_t - y_{t|t-1}) \\ \mathbf{P}_{t|t}^{(xx)} &= \mathbf{P}_{t|t-1}^{(xx)} - K_t \mathbf{P}_{t|t-1}^{(xx)} K_t^T \end{aligned}$$

IV. MULTIPLE OBJECT TRACKING

Following tracking-by-detection methods, which are popular for solving multiple-object tracking tasks, a detector is applied in each frame to generate object candidates which are outputs of the detector. One UKF is adopted for tracking one object separately; thus a group of detected pedestrians defines a family of UKFs to be processed simultaneously. Each UKF tracks one detected object. The predicted state of a UKF is used for data association; when an observation (of the tracked object) is available in the current frame then we update the predicted state by using the corresponding UKF.

A. Detection

Detection-by-tracking methods rely on evaluating rectangular regions of interest, and we call them object boxes if positively identified as containing an object of interest. For pedestrian tracking, we adopt the popular histogram of oriented gradients (HOG) feature method and a support vector machine (SVM) classifier, originally introduced in [19]. HOG features describe the human profile by an oriented gradient histogram. An SVM classifier is able to handle high-dimensional and nonlinear features (such as HOG features). It projects sample features into a high-



Fig. 1. The depth map on top uses a colour code for calculated distances; depth values are only shown at pixels where the mode filter accepts the given value. The lower images show detected (coloured) object boxes.

dimensional space, and then finds a hyperplane to separate two classes. Instead of using a sliding window, regions of interest (i.e., inputs to the classifier) are selected by analysing calculated stereo information (depth and disparity maps), as proposed in [20].

Fig. 1 shows several detection results in pedestrian sequence, dots (cyan) denote the boxes' centre that are recognized as pedestrians, and the red rectangles denote the final detection results. As can be seen in the results, the object boxes may contain background, shift from the object, or miss the pedestrians.

For the detection of *Drosophila* larvae (an example of 2D movement), thresholds and connected components are adopted to obtain one object box for each larva. Several larva detection results are shown in Fig. 2. As the scene is certain, the detection results are more reliable when compared to the pedestrian sequence. However, no depth information is available here.



Fig. 2. Larvae detection results shown by (cyan) object boxes.

B. UKF-based Object Tracking

As there is an unknown number of objects in a scene, the state-dimensionality would expand significantly if we would have decided to track all pedestrians in one UKF; in this case, the speed of tracking reduces dramatically when the scene is crowded with many detected objects. Thus, we decided on one UKF for each detected object for tracking.

Choosing a proper model is important. In this subsection we offer three models for possible selection: *3DVT* means that 3D position (world coordinates) with velocity is observed, *3DT* means that 3D position without velocity is observed, and *2DT* means that 2D position (image coordinates) without velocity is observed. These models are compared in Section V.

1) The Two 3D Models

In the *3DVT* model, the object is tracked in 3D world coordinates. Its 3D position (x, y, z) is the first part of the state. We also include the velocity (v_x, v_y, v_z) . Thus, a state $x=(x, y, z, v_x, v_y, v_z)^T$ is 6-dimensional.

a) Process model: We assume constant velocity between adjacent frames, with Gaussian distributed noise acceleration $n_a \in N(0, \Sigma_{n_a})$. The diagonal elements in Σ_{n_a} are set to be equal and denoted by $\sigma_{n_a}^2$. Thus,

$$\begin{aligned} \dot{x} &= x + (v_x + n_{ax}\Delta t)\Delta t, & \dot{v}_x &= v_x + n_{ax}\Delta t \\ \dot{y} &= y + (v_y + n_{ay}\Delta t)\Delta t, & \dot{v}_y &= v_y + n_{ay}\Delta t \\ \dot{z} &= z + (v_z + n_{az}\Delta t)\Delta t, & \dot{v}_z &= v_z + n_{az}\Delta t \end{aligned}$$

where Δt is the time interval between subsequent frames.

b) Observation model: An observation consists of the position (i_0, j_0) (say, the centroid of the detected object box in the left camera), disparity d of the detected object, and velocity (v_{ox}, v_{oy}, v_{oz}) in 3D coordinates. The usual pinhole camera projection model is used to map 3D points onto the image plane,

$$\begin{aligned} i_o &= fx/z, & j_o &= fy/z \\ d &= fb/z, & v_{ox} &= v_x \\ v_{oy} &= v_y, & v_{oz} &= v_z \end{aligned}$$

where f denotes focal length, and b denotes the length of the baseline between two rectified stereo cameras. In this case,

$$\mathbf{R} = \text{diag}(\sigma_{n_{mp}}^2, \sigma_{n_{mp}}^2, \sigma_{n_{mp}}^2, \sigma_{n_{mv}}^2, \sigma_{n_{mv}}^2, \sigma_{n_{mv}}^2)$$

For the disparity d we select the mode in the disparity map in a fixed (e.g., 20×20) neighbourhood around the centroid (i_0, j_0) of the detected object box. 3D scene flow (v_{ox}, v_{oy}, v_{oz}) can be obtained by combining optic flow and stereo information [21].

As it is difficult to obtain high-quality scene flow as required for *3DVT*, *3DT* simplifies the *3DVT* model by excluding the scene flow from the observation, and has the same process model as *3DVT*. In this case,

$$\mathbf{R} = \text{diag}(\sigma_{n_{mp}}^2, \sigma_{n_{mp}}^2, \sigma_{n_{mp}}^2)$$

2) The 2D Model

If only monocular recording is available, the object is tracked in the 2D image plane only. The state $x=(i, j, v_i, v_j)^T$ consists of position (i, j) and velocity (v_i, v_j) .

a) Process model: The process model is the same as for the 3D models. We assume a constant velocity between subsequent frames with a Gaussian noise distribution for acceleration n_a :

$$\dot{j} = j + (v_j + n_{aj}\Delta t)\Delta t, \quad \dot{v}_j = v_j + n_{aj}\Delta t$$

b) Observation model: An observation consists of the central position (i_0, j_0) of an object box only, $i_0 = i$ and $j_0 = j$,

resulting in $\mathbf{R} = \text{diag}(\sigma_{n_{mp}}^2, \sigma_{n_{mp}}^2)$ for this case.

C. Data Association

As each object is tracked independently, data association by matching candidates to existing trajectories becomes important. If no match is found then we decide to initialize a new tracker.

Since object movements are continuous, the estimated velocity in the UKF can be used as a cue to localize the search area in order to find the match object. For each trajectory, the possible location (i.e., (x_p, y_p, z_p) for 3D, and (i_p, j_p) for 2D) of the object in the current frame is predicted by a process model used in the EKF. This location is used as a reference for searching potentially matching candidates in the current frame. Currently we simply match candidates based on the shortest Euclidean distance and a given threshold T .

One candidate might be matched with several trajectories if the Euclidean distance is below T . Trajectories compete for the candidates, and in the end, the closest one wins. If a candidate is not matched to any trajectory, a new tracker is initialized. If a trajectory does not win any of the candidates, the tracker is propagated with the given prediction, and the new state is the predicted state, without being updated by an observation (because it is not available).

No object appearance description is used here for assigning an object to a trajectory. In general, the inclusion of appearance representation (e.g., a colour histogram, or an instance-specific shape model) improves the performance. However, this is out of the scope of this paper, where we are focusing on the combination of different data association methods.

V. EXPERIMENTS

In this section, first, our three models (3DVT, 3DT, and 2DT) are tested in a simulated environment with different parameter sets. Second, our multiple-object tracking method is tested on real video sequences where (3D example) pedestrians are walking in inner-city scenes, or (2D example) larvae are moving on a flat culture dish.

A. Simulated Tracking

The three models defined in Section IV are tested in a simulation environment in OpenGL (SGI, Fremont, CA, USA). A cub is moving on a circular path around a 3D point with constant speed, as shown in Figs. 3 and 4. Acceleration noise n_a with different covariance (e.g., $\sigma_{n_a}^2 = 0.0001, 0.01, 1$), and measurement noise n_m with different covariance (e.g., $\sigma_{n_{mp}}^2 = 10, 50, 100, \sigma_{n_{mv}}^2 = 50, 100, 150$), are used to test and

compare the three models' performance. The simulation environment is different for 2D and 3D models, where for 2D, positions are integral pixel coordinates in the image plane, but for 3D, position coordinates are reals. The radius of the circle in the 3D models is 10, while in the 2D model it is 50. In both environments, measurements are degraded by noise before being sent to the UKF.

Fig. 3 demonstrates the effect of $\sigma_{n_a}^2$, having fixed $\sigma_{n_{mp}}^2$ and $\sigma_{n_{mv}}^2$. Experiments show that larger $\sigma_{n_a}^2$ values result in more unstable trajectories. A large $\sigma_{n_a}^2$ means that the process model produces a predicted state that is fluctuating with large magnitudes. Results show that $\sigma_{n_a}^2 = 0.0001$ is a reasonable choice for 3D models. For the 2D model, a smaller $\sigma_{n_a}^2$ yields smooth estimation, but the shifts are significant.




Fig. 3. Simulation results for variations in the variance of acceleration noise. From left to right, $\sigma_{n_a}^2 = 0.0001, 0.01, \text{ or } 1$, with fixed values $\sigma_{n_m}^2 = \sigma_{n_p}^2 = 50$, and $\sigma_{n_v}^2 = 100$. From top to down, the tracking model is 3DVT, 3DT and 2DT, respectively.

A larger $\sigma_{n_a}^2$ value produces estimations that are closer to the true positions, but fluctuations are significant, for an experiment with $\sigma_{n_a}^2 = 1$ for the 2D case. In general, 3DT and 3DVT converge better than 2DT, while 3DT and 3DVT show a similar performance. 3D models use stereo information rather than just a single image as for the 2D model, which also proves that stereo information can help to improve the tracking performance. As the measured 3D position is noisy, the measure of velocity is even noisier; this appears to be the main reason for the observation that the inclusion of velocity cannot improve the performance.

Fig. 4 shows results for our models for different covariance values $\sigma_{n_{mp}}^2$ and $\sigma_{n_{mv}}^2$ of measurement noise. Significantly increasing measurement noise (i.e., higher uncertainty of observations) reduces the performance only slightly. This demonstrates that, to some degree, the UKF is a robust tracker, which is not vulnerable to detection uncertainties. As before, 3DT and 3DVT converge better than 2DT, while 3DT and 3DVT show a similar performance.




Fig. 4. Simulation results for variable variance of measurement noise. From left to right, $\sigma_{n_{mp}}^2 = 10, 50, \text{ or } 100$, $\sigma_{n_{mv}}^2 = 50, 100, 150$, respectively, with fixed $\sigma_{n_a}^2 = 0.0001$ for 3D models, $\sigma_{n_a}^2 = 1$ for 2D models. From top to bottom, the tracking models are 3DVT, 3DT and 2DT, respectively.

B. Multiple Object Tracking in Real Data

In this section we report on the performance of UKF-supported tracking for multiple larvae using the 2DT model, and for multiple pedestrians in traffic scenes using the 3DT model. The larvae and pedestrian sequences are recorded at 30 and 15 frames per second, respectively.

Results for larvae tracking are shown in Fig. 5. As the velocity in the model is initialized by (0,0), the UKF-estimation is “slower” than the real speed of the larvae in the first 30 frames. The speed of convergence can be improved by increasing σ_{n_a} , but it should be noted that the larger the σ_{n_a} value is, the larger the magnitude of fluctuation. The estimated trajectories follow the moving larvae effectively, mainly because all of the larvae are properly detected in all of the frames. However, such a complete detection cannot be expected for pedestrian sequences. Next, we test the UKF for such “noisy” detection results as pedestrian sequences.

The results for pedestrian tracking are shown in Fig. 6. Objects are missing or shifting from time to time due to the clustered background (e.g., the car in the traffic scene detected as a pedestrian), illumination variations leaving some pedestrians undetected, or internal variations between objects (i.e., unstable detections). Our experiments verified that UKF predictions are able to follow irregularly moving pedestrians when detection fails for a few frames, and can even correct unstable detections.

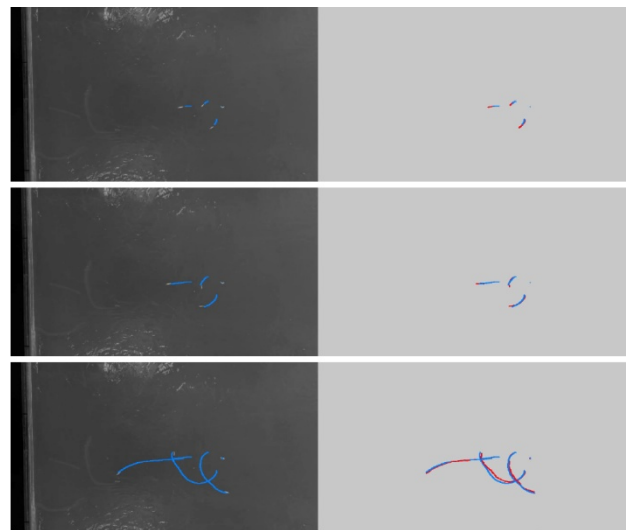


Fig. 5. 2D Tracking results of larva sequences. From top to bottom: tracking results in Frames 26, 46, and 166 of one sequence. The red lines show the detected track, and the white lines show the unscented Kalman filter-predicted track. The blue lines represent estimated trajectories. The left column is the original intensity image overlaid with the estimated trajectories.

The second frame in Fig. 6 shows that the undetected pedestrian is predicted correctly in the white object box and is successfully matched to a detected position in the third frame. The last frame in Fig. 6 demonstrates that displaced detections are corrected by the UKF. Using only the defined distance rule for data assignment, this appears to be insufficient, especially for the given detection results. A small threshold may lead to a mismatch (i.e., the detection fails to satisfy the rule), and a large threshold may lead to an identity switch (i.e., a pedestrian is matched to another pedestrian).

VI. CONCLUSIONS

Assigning one UKF to each detected (moving) object simplifies the design and implementation of UKF prediction of 2D or 3D motion. Experiments demonstrate the robustness of the chosen approach. This tracker only generates short-term tracks when detection is not reliable; long-term tracking should be possible by also introducing dynamic programming.

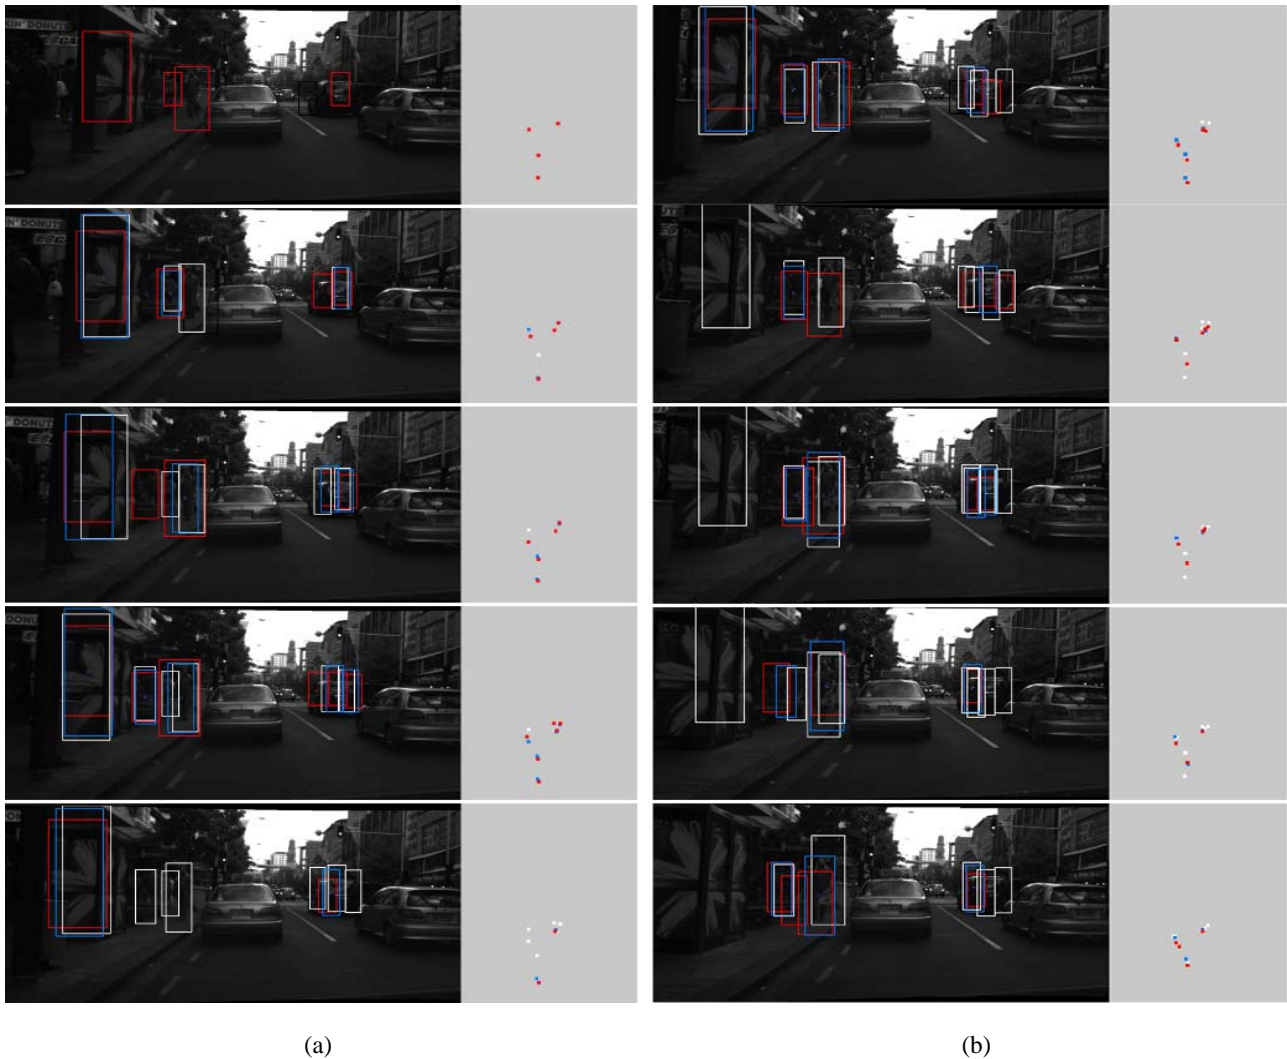


Fig. 6. (a) Object boxes (red), unscented Kalman filter (UKF) predictions (white), and UKF estimations (blue) are overlaid on the original intensity image. The black box is an invalid detection excluded by a disparity check. (b) coloured blocks denote the object boxes (red), UKF predictions (white), and UKF estimations (blue) in the XZ-plane in real-world coordinates.

For evaluating the performance in real-world (either 2D or 3D) applications, more extensive tests need to be undertaken, especially for the design and evaluation of quantitative performance measures. For example, the measures discussed in [22] for evaluating visual odometry techniques might also be of relevance for the tracking case.

ACKNOWLEDGMENTS

The authors thank *Simon Hermann* for providing his implementation of semi-global matching for stereo analysis, *Gabriel Hartmann* (both Auckland) for support regarding the unscented Kalman filter, and *Benjamin Risse* (Münster) for video data with moving *Drosophila* larvae.

REFERENCES

- [1] B. Wu and R. Nevatia, "Detection and tracking of multiple, partially occluded humans by Bayesian combination of edgelet based part detectors," *International Journal of Computer Vision*, vol. 75, no. 2, pp. 247-266, 2007.
- [2] M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. van Gool, "Online multiperson tracking-by-detection from a single, uncalibrated camera," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 9, pp. 1820-1833, 2011.
- [3] A. Ess, K. Schindler, B. Leibe, and L. van Gool, "Object detection and tracking for autonomous navigation in dynamic environments," *International Journal of Robotics Research*, vol. 29, 14, pp. 1707-1725, 2010.

- [4] B. Leibe, K. Schindler, N. Cornelis, and L. van Gool, "Coupled object detection and tracking from static cameras and moving vehicles," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 10, pp. 1683-1698, 2008.
- [5] M. Meuter, U. Iurgel, S. B. Park, and A. Kummert, "The unscented Kalman filter for pedestrian tracking from a moving host," *Proceedings of IEEE Intelligent Vehicles Symposium*, Eindhoven, Netherlands, pp. 37-42, 2008.
- [6] E. A. Wan and R. van der Merwe, "The unscented Kalman filter for nonlinear estimation," *Proceedings of IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium*, Lake Louise, AB, pp. 153-158, 2000.
- [7] A. Yilmaz, O. Javed, and M. Shah, "Object tracking: a survey," *ACM Computing Surveys*, vol. 38, no. 4, article no. 13, 2006.
- [8] A. Ess, B. Leibe, K. Schindler, and L. van Gool, "Robust multiperson tracking from a mobile platform," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 10, pp. 1831-1846, 2009.
- [9] W. F. Leven and A. D. Lanterman, "Unscented Kalman filters for multiple target tracking with symmetric measurement equations," *IEEE Transactions on Automatic Control*, vol. 54, no. 2, pp. 370-375, 2009.
- [10] D. Mitzel, E. Horbert, A. Ess, and B. Leibe, "Multi-person tracking with sparse detection and continuous segmentation," *Computer Vision ECCV 2010, Lecture Notes in Computer Science*, vol. 6311, pp. 397-410, 2010.
- [11] D. Mitzel, P. Sudowe, and B. Leibe, "Real-time multi-person tracking with time-constrained detection," *Proceedings of the British Machine Vision Conference*, Dundee, UK, pp. 104.1-104.11, 2011.
- [12] D. Mitzel and B. Leibe, "Real-time multi-person tracking with detector assisted structure propagation," *Proceedings of Computational Methods for the Innovative Design of Electrical Devices*, Barcelona, Spain, pp. 974-981, 2011.
- [13] M. M. Shaikh, B. Wook, C. Lee, T. Kim, T. Lee, K. Kim, and D. Cho, "Mobile robot vision tracking system using unscented Kalman filter," *Proceedings of IEEE/SICE International Symposium on System Integration*, Kyoto, Japan, pp. 1214-1219, 2011.
- [14] G. Welch and G. Bishop, "An introduction to the Kalman filter," University of North Carolina at Chapel Hill, Chapel Hill: NC, *TR 95-041*, 1995.
- [15] S. Huang and G. Dissanayake, "Convergence analysis for extended Kalman filter based SLAM," *Proceedings of IEEE International Conference on Robotics and Automation*, Orlando, FL, pp. 412-417, 2006.
- [16] H. S. Yazdi and S. E. Hosseini, "Pedestrian tracking using single camera with new extended Kalman filter," *International Journal of Intelligent Computing and Cybernetics*, vol. 1, no. 3, pp. 379-397, 2008.
- [17] Y. Cai, N. de Freitas, and J. J. Little, "Robust visual tracking for multiple targets," *Proceedings of the 9th European Conference on Computer Vision*, Graz, Austria, pp. 107-118, 2006.
- [18] G. Hartmann, "Unscented Kalman filter sensor fusion for monocular camera localization," MSc thesis, University of Auckland, New Zealand, 2012.
- [19] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, San Diego, CA, pp. 886-893, 2005.
- [20] D. M. Gavrilu and S. Munder, "Multi-cue pedestrian detection and tracking from a moving vehicle," *International Journal of Computer Vision*, vol. 73, no. 1, pp. 41-59, 2007.
- [21] A. Wedel, T. Brox, T. Vaudrey, C. Rabe, U. Franke, and D. Cremers, "Stereoscopic scene flow computation for 3D motion understanding," *International Journal of Computer Vision*, vol. 95, no. 1, pp. 29-51, 2011.
- [22] R. Jiang, R. Klette, and S. Wang, "Statistical modeling of long-range drift in visual odometry," *Computer Vision ACCV 2010 Workshops, Lecture Notes in Computer Science*, vol. 6469, pp. 214-224, 2011.



Junli Tao

Ms. Tao is a PhD student at Auckland University, New Zealand. She received her Master's degree in 2010 from Hunan University in Hunan, China. Her research interests are in active pedestrian protection systems for diver assistance, described by keywords such as ego-motion analysis, pedestrian detection and tracking, and pedestrian intention prediction.



Dr. Reinhard Klette

Dr. Klette is a professor in the Computer Science Department at Auckland University, New Zealand. His professional interests are in computer vision (in theory and applications) and in the design of geometric algorithms. He has (co-) authored more than 350 peer-reviewed publications. He is the Editor-in-Chief of the *Journal on Control Engineering and Technology*.