

Application of the Particle Filter to Tracking of Fish in Aquaculture Research

Tomasz Pinkiewicz, Ray Williams
School of Computing and Information Systems
University of Tasmania
Tasmania, Australia
tomaszp@utas.edu.au

John Purser
National Centre for Marine Conservation
and Resource Sustainability
AMC, University of Tasmania
Tasmania, Australia

Abstract

The analysis of fish movement as an indicator of fish behaviour plays an important role in aquaculture research. Currently observations are carried out manually using video recordings. In this paper we describe a tracking system which can automatically detect and track two fish in a video sequence in a small aquaculture tank. The system is based on the particle filter tracking algorithm augmented by an adaptive partition scheme and using a Global Nearest Neighbour approach for data association. Results show that this method is sufficient for simple interactions where fish bypass each other without significant changes in velocity. However, more complex scenarios involving occlusions, loss of tracks and fish manoeuvres can cause ambiguity during data association.

1. Introduction

Fish show a range of behaviours and responses to factors present within the captive environment, be it in sea cages or small tanks [1]. By way of example: fish may display sudden aggressive movements towards other fish when competing for limited food, swim at different speeds and depths relative to the phase of the meal, reduce swimming speed and increase response times when moribund (sick), display leaping behaviour when infested with parasites and show a flight response when attacked by predators such as seals. One factor common to many of these situations is that the behaviour is based on movement with speed and location as co-factors. Hence approaches that could be used to measure such behaviour include tracking, spatial distribution, activity rhythms and activity levels. In small tanks, with low numbers of fish, the emphasis is on individual behaviours and on spatiotemporal distributions. Analysis of footage recorded over a 12 hour duration can take up to 2-3 days using manual observation. This research attempts to automate the detection and tracking of small number of fish in

small tanks with the resultant movement data being used for more complex behavioural analysis.

1.1. Use of video analysis in aquaculture

The use of video in aquaculture is widespread both in research and commercial facilities but most of the analysis has been done manually and on an ad hoc basis to answer specific questions. This widespread use is due to its low cost and non invasiveness compared to other methods (e.g. tagging). However the quality of data can vary as videos can often be blurred due to motion and compression, and light reflections in the water can cause loss of data. While steps are taken to eliminate reflections, one has to be mindful not to alter the environment in a way that will affect the fish. For this reason, the use of lateral illumination and clear tanks is not possible.

Automated analysis of images or video has mainly been focused on the determination of fish biomass using stereo camera systems [2]. More recent efforts have shifted to the detection and identification of fish. Lee et al. [3] used motion and fish contours to detect and identify wild fish in river passageways. Williams et al. [4] used image processing and an Active Shape Model to detect captive fish in sea cages as part of a wider study to detect marine life in underwater images but so far there has been no attempt to automatically analyse fish movement in tanks or sea cages for the purpose of behavioural analysis. This research is progressing towards the automated analysis of a small numbers of fish in small research tanks. There are two major objectives for the project: to track individual fish for long periods of time (10-12 hours per day) and to detect fish presence for the purpose of spatiotemporal analysis. The latter task does not require long term tracking or identification; rather we're interested in the areas fish occupy and for how long. A secondary aim is to provide velocity profiles of fish responding to feeding or changes in the environment throughout the day. Again long term tracking is not required here; only short term tracks are necessary to determine the velocity. This paper

describes the basic framework to achieve these aims at least for short video sequences. We start by tracking two fish but the final goal of this research is to track 4-6 fish, so we can learn about fish hierarchy and how it affects feeding, growth and welfare.

1.2. Background on particle filters

A particle filter is a sequential estimation method which uses a sampling based approach to estimate the state of a stochastic dynamic system. In our case the system is the movement of fish in a tank, observed using video, and the state is a set of variables describing fish positions and velocities represented as state vectors. Each particle is a random sample containing a state vector. A set of particles represents the state of the system through calculation of a *posterior distribution* which can be described using mean and covariance (other methods of describing the distribution exist). The method is derived from Monte Carlo methods and is also known as the Sequential Monte Carlo method (SMC). The filter works in two stages: the prediction stage and the update stage. In the prediction stage, particles at time k are evolved from particles at time $k - 1$ using a dynamic model combined with the *posterior density* from time step $k - 1$ to produce a set of particles representing the *predicted density*. The update stage modifies the *predicted density* with measurements available at time k by calculating weights for all particles (the *observation density*) and constructs the *posterior density*. To prevent degeneration of particles over time, a resampling step is introduced which removes particles with low weights and splits particles with high weights. Once the *posterior density* at time k is known, the point estimate is obtained by calculating the mean or *maximum a posteriori* (MAP) of this density. For a more detailed explanation of particle filter concepts see [5, 6].

Particle filters have become popular over the last decade due to their ability to estimate non-linear systems (albeit only through approximation). Early work was undertaken in the signal processing field [7] but soon the computer vision community recognised its potential for detection and tracking of objects in video sequences. One of the most widely known early work in the computer vision field was the development of the CONDENSATION algorithm [8].

Initial versions of the particle filter suffered from one limitation: they were unable to maintain multi-modal state densities when tracking multiple targets. Since then much research has been carried out to address this issue. McCormick and Blake [9] proposed a probabilistic exclusion principle while Schulz et al. [10] incorporated the Joint Probabilistic Data Association Filter (JPDAF) directly into the particle filter to allow multi-target tracking. Another approach was to use a mixture of particle filters and allow them to interact during the computation of weights [11].

In computer vision particle filters often take advantage of the colour information available in the scene to track multiple, similar targets [12]. Various other approaches have been proposed. They include the Kernel Particle Filter [13] and the more recent hybrid colour particle filter [14] which originates from the track-before-detect concept [6]. The former enhances the particle filter with a mean-shift approach to enable tracking of like targets when in close proximity and during occlusion. The latter uses a discrete variable to denote the number of targets each particle is tracking. Also a colour based particle filter has been combined with an Adaboost classifier to improve the tracking of ice hockey players and to deal with the movement of the camera and the frequent appearance and disappearance of players [15].

Another improvement to the particle filter is called the Adaptive Partition scheme [16]. This scheme adds a bias to the particle importance weighting and allows the particle filter to utilise current measurements during particle prediction. The result is that the particle filter can track multiple targets more reliably with fewer particles used.

1.3. Paper outline

In this paper we demonstrate how a hybrid colour particle filter can be used to track fish in small aquaculture tanks in the presence of occlusions. Section 2 will describe the particle filter and the modifications we carried out to suit our application. Section 3 will show how the particle filter has been augmented with an Adaptive Partition approach and will outline a basic data association scheme, based on the Global Nearest Neighbour technique, which utilises the output of the particle filter to maintain target identification. In Section 4, we'll demonstrate the outcomes of the tracking system when dealing with two fish. Finally, we will conclude in Section 5 with a summary and discussion of the results and further research directions for this project.

2. Colour particle filter

Our detection system is based on the hybrid colour particle filter by Czyz et al. [14], which assumes that the number of targets is known. This algorithm has been developed to deal with tracking of multiple like objects in video using colour information contained in the video sequence. We define the vector y_k to be:

$$y_k = \begin{cases} E_k & \text{if } E_k = 0, \\ [x_{1,k}^T, E_k]^T & \text{if } E_k = 1, \\ [x_{1,k}^T, x_{2,k}^T, E_k]^T & \text{if } E_k = 2, \\ \vdots & \vdots \\ [x_{1,k}^T, \dots, x_{M,k}^T, E_k]^T & \text{if } E_k = M, \end{cases} \quad (1)$$

where y_k is a combined state vector, $x_{l,k}$ is a state vector of object $l = 1 \dots E_k$ at time k and E_k is the discrete existence variable which signifies the number of targets existing in the state vector y_k . This arrangement means that different particles of the filter may have different numbers of state vectors $x_{l,k}$ depending on the value of E_k . The objective is to find the posterior density $p(y_k|Z_k)$, where $Z_k = \{z_i, i = 1 \dots k\}$, z_i being an image frame at time i . We find this density using Bayesian estimation in two steps. The prediction step computes the *predicted density* using available knowledge from the previous time step and the dynamic model which evolves the state. It is represented as:

$$p(y_k|Z_{k-1}) = \int p(y_k|y_{k-1})p(y_{k-1}|Z_{k-1})dx_{k-1}, \quad (2)$$

where $p(y_k|y_{k-1})$ is the *transitional density* which defines the evolution of the stochastic dynamic state and $p(y_{k-1}|Z_{k-1})$ is the *posterior density* at the previous time step. The *posterior density* $p(y_k|Z_k)$ is then calculated in the update stage:

$$p(y_k|Z_k) \propto p(z_k|y_k)p(y_k|Z_{k-1}), \quad (3)$$

where $p(z_k|y_k)$ is an *observation density* also known as the likelihood function in Bayes' Theorem.

The existence variable E_k is propagated from E_{k-1} using a Markov chain specified by the *transitional probability matrix* (TPM). In our case the TPM for 2 targets is:

$$\Pi = \begin{bmatrix} (1 - P_b) & P_b & 0 \\ P_d & (1 - P_d - P_m) & P_m \\ 0 & P_r & (1 - P_r) \end{bmatrix}, \quad (4)$$

where P_b, P_d are probabilities of object "birth", "death" respectively and P_m, P_r are probabilities that the number of targets will increase or decrease respectively. For our application we empirically choose these values to be $P_b = 0.95$, $P_d = 0.05$, $P_m = 0.15$, $P_r = 0.5$. High probability of birth means that the first fish will be acquired quickly, while low probability of death is because the fish is unlikely to exit the area (due to constraints of the tank). The value we determined for P_m allows to detect the second fish quickly, while higher P_r allows to quickly cease tracking of the second fish if the detection is not possible due to reflections on the water surface. The propagation is carried out using a regime transition algorithm from [6] (Table 3.9).

2.1. Extraction of measurements

The likelihood function of the filter requires comparison between the reference histogram and the histogram computed from a region within the current image z_k . We used the Hue Saturation Value (HSV) colour space to produce histograms because it is more robust to illumination

changes than the RGB colour space. The HSV histogram consists of $N = N_h * N_s + N_v$ bins [17] and in our experiments $N_h = N_s = N_v = 7$. The reference histogram was generated from a mean of 20 reference images of fish in various areas of the tank and different illumination levels. To get the distance between the reference histogram and the region histogram $D_k = dist[q^*, q_k]$, we used the Bhattacharyya similarity coefficient which is defined as:

$$D_k = \sqrt{1 - \sum_{u=1}^U \sqrt{q^*(u)q_k(u)}}, \quad (5)$$

where U is the number of bins, q^* is the reference histogram, q_k is the compared region in image z_k . The likelihood function for multi-object observations is calculated thus:

$$p(z_k|Y_k^m, E_k = m) \propto \frac{1}{\sqrt{2\pi\sigma}} \exp\left\{-\frac{1}{2\sigma^2} \sum_{i=1}^m D_{i,k}^2\right\}, \quad (6)$$

where m is a number of objects in the area and i refers to the i th object within the state vector $Y_k^m \equiv [y_k^1, \dots, y_k^m]$. Therefore the histogram q_k is build from the region defined by $x_{i,k}$ within the image z_k .

3. Implementation

We now consider a particle filter as a set of N random samples. The state vector y_k can be rewritten to include an additional index:

$$y_k^n = [E_k^n, x_{1,k}^n, \dots, x_{g,k}^n] \quad (n = 1, \dots, N, g = E_k^n)$$

The likelihood for particle n can be calculated using Eq. (6) but we also have to consider the impact of the background on the detection process. So the likelihood function is

$$L_k^n(E_k^n) = \exp\left\{-\frac{1}{2\sigma^2} \sum_{i=1}^{E_k^n} [(D_{i,k}^n)^2 - (D_{i,k}^{n,B})^2]\right\}, \quad (7)$$

where $D_{i,k}^{n,B}$ is the distance between the reference histogram and the histogram of the background at coordinates specified by $x_{i,k}^n$. Czyz et al. [14] used a constant for background estimation because of the movement of the camera but our experiments have a static camera, therefore we estimated the colour background using a hybrid background maintenance algorithm [18]. The un-normalised importance weights are calculated based on the value of E_k^n :

$$\tilde{w}_k^n = \begin{cases} 1 & \text{if } E_k^n = 0, \\ L_k^n(E_k^n) & \text{if } E_k^n > 0. \end{cases} \quad (8)$$

3.1. The adaptive partition method

The use of the *transitional density* $p(y_k|y_{k-1})$ to propose new particles can be inefficient because it requires a large number of particles especially when tracking multiple targets. Kreucher et al. [16] proposed an adaptive partition method which improves the proposal by incorporating current measurements. While calculation of individual particles requires more computational effort, fewer particles are required. The adaptive partition method consists of two submethods.

The first submethod is the Independent Partition (IP) method and is used when targets are well separated. The basic principle behind this submethod is that for every target, the method samples a new set of particles from the *transitional density* and computes the weights of each particle. After normalising weights, a new set of particles is drawn according to the distribution of weights. This new set of particles becomes the prediction density. In addition, a bias is introduced, which is used when calculating global weights.

The second submethod is the Coupled Partition (CP) method and is used when targets are close together or occluding. For every particle of every target, the method samples R realisations of the corresponding state vector within each particle from the transitional density and computes weights for each realisation. An index is then drawn based on the distribution of weights and a new particle is selected from realisations based on the index. Again a bias is introduced for the global weight calculation to improve the selection process.

Because the Independent Partition method requires that the state vector (partition) ordering is consistent in all particles (i.e. state vectors $x_{1,k}^i$ and $x_{1,k}^j$, where $i \neq j$ belong to the same targets), an additional step called Partition Sorting needs to be carried out. This step uses the K-means algorithm to provide a proper ordering of partitions.

3.2. The revised particle filter algorithm

We have incorporated the adaptive partition scheme into the hybrid colour particle filter to deal better with tracking multiple targets. Table 1 outlines the modified algorithm for N number of particles. The particles are still drawn in Step 2a(i) and 2b(i) based on the pair of existence variables (E_{k-1}^n, E_k^n) . Czyz et al. [14] describes this selection procedure and how it controls the appearance and disappearance of targets.

An additional modification was to the likelihood function. To prevent the appearance of objects on the same image region Czyz et al. [14] included a condition which checked the distance between the state vector within a particle and set the weight of the particle to zero if the distance was below a given threshold. We decided to retain that but

Table 1. Particle Filter pseudo-code

$[\{y_k^n\}_{n=1}^N] = \mathbf{PF}[\{y_{k-1}^n\}_{n=1}^N, z_k]$

- (1) $[\{E_k^n\}_{n=1}^N] = \mathit{ETrans}[\{E_{k-1}^n\}_{n=1}^N, \Pi]$
- (2) FOR $t = 1 : T_{max}$
 - a. IF $\min_{j \neq t} \|\hat{x}_{t,k-1} - \hat{x}_{j,k-1}\| > \tau$
 - a(i) $[\{x_{t,k}^n, b_t^n\}_{n=1}^N] = \mathit{IP}[\{x_{t,k}^n\}_{n=1}^N, z_k]$
 - b. ELSE
 - b(i) $[\{x_{t,k}^n, b_t^n\}_{n=1}^N] = \mathit{CP}[\{x_{t,k}^n\}_{n=1}^N, z_k]$
 - c. ENDIF
 - (3) END FOR
 - (4) FOR $n = 1 : N$
 - a. Evaluate importance weight \tilde{w}_k^n
 - b. Apply bias $\tilde{w}_k^n = \frac{\tilde{w}_k^n}{\prod_{t=1}^{E_k^n} b_{n,t}}$
 - (5) END FOR
 - (6) FOR $n = 1 : N$
 - (7) Normalise weights: $w_k^n = \frac{\tilde{w}_k^n}{\sum_{m=1}^N \tilde{w}_k^m}$
 - (8) END FOR
 - (9) Resample: $[\{y_k^n\}_{n=1}^N] = \mathit{RESAMPLE}[\{y_k^n, w_k^n\}_{n=1}^N]$
 - (8) Partition Sorting of $\{y_k^n\}_{n=1}^N$
 - (9) Calculate number of targets T and estimates $\{\hat{x}_{t,k}\}_{t=1}^T$

only if a particle increased the number of targets it contained (i.e. a new state vector was added to the particle). For the remainder of the particles the distance check was included into the global likelihood function (Table 1, Step 4a) which would favour particles which are further away. Eq. (9) computes the minimum distance between state vectors within the particles and the result is used in Eq. (10) to create a new likelihood function.

$$d(y_k^n) = \min_{i \neq j} \left\{ \left[\left(\|x_{i,k}^n - x_{j,k}^n\| \right)_{i=1}^{E_k^n} \right]_{j=1}^{E_k^n} \right\} \quad (9)$$

$$\tilde{L}_k^n(E_k^n) = L_k^n(E_k^n) \exp \left\{ -H d(y_k^n) \right\} \quad (10)$$

$H = 1000$ in our experiments and this will depend on the nature of the problem. This solution produces acceptable results when targets are separated but when targets are close and the Coupled Partition method is activated, we reverted to the original likelihood in Eq. (7) for 3 time steps. This was enough time for the Coupled Partition method to propagate particles based on the kinematic motion of targets. After that Eq. (10) was re-applied to allow de-coupling of particles. The likelihood function used within the Independent Partition and the Coupled Partition method (Table 1, Step 2a(i) and 2b(i)) was Eq. (7) because the state vectors were considered independently.

3.3. Data Association

The number of targets, target estimates and variances were calculated using equations from Section III E in [16]. Estimates with high variance were discarded and not used in data association. For data association we used the Global Nearest Neighbour method. The method works by solving the two dimensional assignment problem using Munkres' algorithm [19]. Input to the algorithm is an assignment matrix which contains the squared Euclidean distance d^2 from each target to each observation. d^2 is defined in the following way:

$$d^2 = (x_t - x_m)^2 + (y_t - y_m)^2 + \alpha((\dot{x}_t - \dot{x}_m)^2 + (\dot{y}_t - \dot{y}_m)^2), \quad (11)$$

where indexes t and m refer to target and measurement respectively, x, y are centroid co-ordinates, \dot{x}, \dot{y} are velocity components. Extra weighting is given to the velocity using α . In our experiments $\alpha = 10$. This heavy weighting helps to resolve the ambiguity in the data association when fish are close but move in different directions and at different speeds.

4. Results

We analysed a 59 second video sequence of two fish in a small aquaculture tank. The video was recorded at 29.5 frames per second (fps). This gave a total of 1742 frames but we only examined every 7th frame, so the practical frame rate was 4.2fps and 249 frames were processed. There were three major interactions between the fish during this sequence. Each rectangle in the following figures represents a 160×120 image frame. Example image frames are shown in Fig. 3, 4 and 5.

Figure 1 shows the first interaction in which Fish 1 approaches Fish 2 and suddenly changes speed and direction and swims away from Fish 2. In frames 12, 13 and 14 both fish approach each other. Frames 15 and 16 show contact and the beginning of the manoeuvre by Fish 1. Estimates in frame 16 are ambiguous and this is an effect of the Coupled Partition method. Frame 17 does not produce any estimates due to high variance. Estimates in frame 18 are correctly matched due to the use of weighted velocity in the data association stage (Section 3.3). Velocity plays an important role in this scenario because Fish 1 moves much faster than Fish 2 and this characteristic helps to differentiate them. Finally frame 19 shows the two fish well separated.

The second scenario (Fig. 2) shows the fish passing each other in very close proximity. Fish 1 is above Fish 2 in 3D but the 2D perspective creates an occlusion. The fish join in frame 32 and continue to bypass each other until frame 37. During this time the Couple Partition method is active and

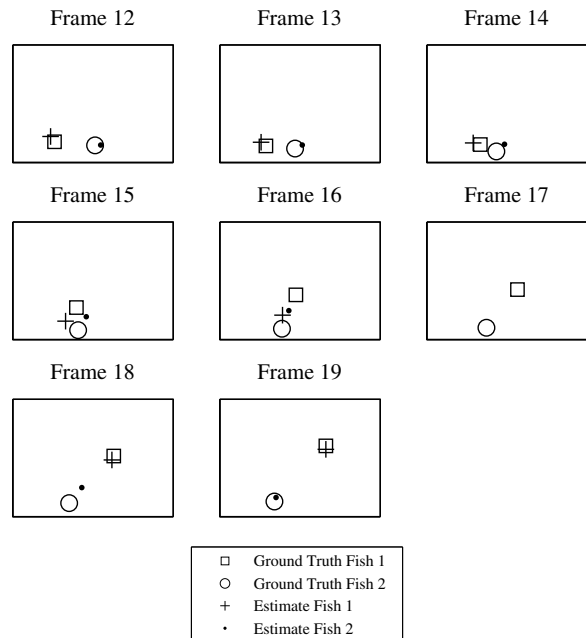


Figure 1. Aggressive fish interaction.

frame 36 shows that both estimates show the same target. But in frame 37, the modified likelihood (Eq. 10) allows for the particle filter to be decoupled.

The last scenario (Fig. 6) is the longest and most difficult as it involves a complex manoeuvre by Fish 1 and the inability of the system to track fish due to light reflection on the water surface and close proximity of the two fish. The fish approach each other in frame 57 and Fish 1 carries out a change-of-direction manoeuvre from frame 59 to frame 68. As the fish start separating from frame 68 onwards the lack of detections due to reflections and a relatively parallel movement of the fish causes problems with the data association which becomes evident in frame 80 where Estimate 2 is assigned to Fish 1. This continues until frame 84 when the tracker re-detects Fish 2.

5. Conclusions

The tracking system performs well when the motion of the fish doesn't change significantly during the occlusion (Fig. 2). In other cases the system recovers under changes in velocity but robustness is a key question which hasn't been investigated yet. While Fig. 6 showed that the tracker correctly associated measurements with fish under heavy occlusions and abrupt change in motion, repetition of this experiment showed that the tracker accuracy varied. This might suggest that the system will fail to associate properly using GNN method in more severe situations.

The Coupled Partition mode still requires more investi-

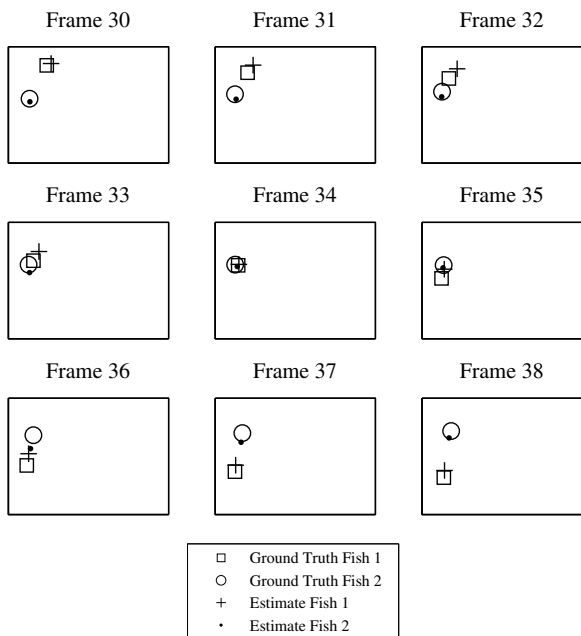


Figure 2. Fish bypassing with occlusion.

gation so that particles are propagated based on their kinematic motion for longer. This will help to avoid situations such as in frame 16 in Fig. 1, where both estimates are assigned to the same region. This suggests that the particle filter failed in its multi-target tracking and deteriorated to a single target mode. While the use of Eq. 10 after 3 time steps helped to separate targets once again (frame 18 in Fig. 1 and frame 84 in Fig. 6) and return to a multimodal particle filter, further work is required to make this system more robust. The use of the Independent Partition method allows for better response to fish manoeuvres with a decreased number of particles.

Further research will focus on multi-dimensional assignment (MDA) data association instead of the GNN method and improvements to the Coupled Partition method. Additional improvement to the particle filter may be achieved using motion [20] or shape data. These improvements will allow us to upscale experiments to larger numbers of fish.

Acknowledgements

This research was partially supported by the Institutional Research Grant P0015639 from University of Tasmania. Authors would like to acknowledge the assistance of the technical staff: Tony Gray, Detlef Planko and Andrew Spilling.

References

- [1] E. Brännäs, A. Alanärä, and C. Magnhagen, “The social behaviour of fish,” in *Social Behaviour in Farm Animals*, L. Keeling and H. Gonyou, Eds. CABI Publishing, 2001, pp. 275–304.
- [2] A. Naiberg and J. J. Little, “A unified recognition and stereo vision system for size assessment of fish,” in *Proceedings of the 2nd IEEE Workshop on Applications of Computer Vision, 1994*, 1994, pp. 2–9.
- [3] D. J. Lee, S. Redd, R. Schoenberger, X. Xiaojian, and Z. Pengcheng, “An automated fish species classification and migration monitoring system,” in *Proceedings of the 29th Annual Conference of the IEEE*, 2003, vol. 2, pp. 1080–1085.
- [4] R. Williams, T. Lambert, A. Kelsall, and T. Pauly, “Detecting marine animals in underwater video: Let’s start with salmon,” in *Proceedings of 12th Americas Conference on Information Systems*, Mexico, Acapulco, 2006.
- [5] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo methods in practice*, ser. Statistics for engineering and information science. New York: Springer, 2001.
- [6] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA: Artech House, 2004.
- [7] N. Gordon, D. Salmond, and A. Smith, “Novel approach to nonlinear/non-gaussian bayesian state estimation,” *IEE Proceedings F: Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, 1993.
- [8] M. Isard and A. Blake, “Condensation – conditional density propagation for visual tracking,” *International Journal of Computer Vision*, vol. 29, no. 1, pp. 5–28, 1998.
- [9] J. MacCormick and A. Blake, “A probabilistic exclusion principle for tracking multiple objects,” in *Proceedings of International Conference on Computer Vision*, 1999, vol. 1, pp. 572–578.
- [10] D. Schulz, W. Burgard, D. Fox, and A. B. Cremers, “Tracking multiple moving objects with a mobile robot,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, Kauai, Hawaii, 2001.
- [11] J. Vermaak, A. Doucet, and P. Pérez, “Maintaining multimodality through mixture tracking,” in *Proceedings 9th IEEE International Conference on Computer Vision, 2003*, 2003, vol. 2, pp. 1110–1116.

- [12] P. Pérez, C. Hue, J. Vermaak, and M. Gangnet, “Color-based probabilistic tracking,” in *Proceedings of the 7th European Conference on Computer Vision-Part I*. Springer-Verlag, 2002, vol. 2350, pp. 661–675.
- [13] C. Chang, R. Ansari, and A. Khokhar, “Multiple object tracking with kernel particle filter,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2005)*, R. Ansari, Ed., vol. 1, 2005, pp. 566–573.
- [14] J. Czyz, B. Ristic, and B. Macq, “A particle filter for joint detection and tracking of color objects,” *Image and Vision Computing*, vol. 25, no. 8, pp. 1271–1281, 2007.
- [15] K. Okuma, A. Taleghani, N. d. Freitas, J. J. Little, and D. G. Lowe, “A boosted particle filter: Multitarget detection and tracking,” in *Proceedings of European Conference on Computer Vision (ECCV04)*, 2004, pp. 28–39.
- [16] C. Kreucher, K. Kastella, and A. O. Hero, “Multitarget tracking using the joint multitarget probability density,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396–1414, 2005.
- [17] Y. Cai, “Robust visual tracking for multiple targets,” Master Thesis, University of British Columbia, 2005.
- [18] S.-Y. Chien, Y.-W. Huang, B.-Y. Hsieh, S.-Y. Ma, and L.-G. Chen, “Fast video segmentation algorithm with shadow cancellation, global motion compensation, and adaptive threshold techniques,” *IEEE Transactions on Multimedia*, vol. 6, no. 5, pp. 732–748, 2004.
- [19] J. Munkres, “Algorithms for the assignment and transportation problems,” *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.
- [20] P. Pérez, J. Vermaak, and A. Blake, “Data fusion for visual tracking with particles,” *Proceedings of the IEEE*, vol. 92, no. 3, pp. 495 – 513, 2004.



Figure 3. Fish begins a manoeuvre. The situation described in Fig.1



Figure 4. Fish bypassing each other. The situation described in Fig.2



Figure 5. Light reflection on the water surface decreases the performance of the tracker. The situation described in Fig.6

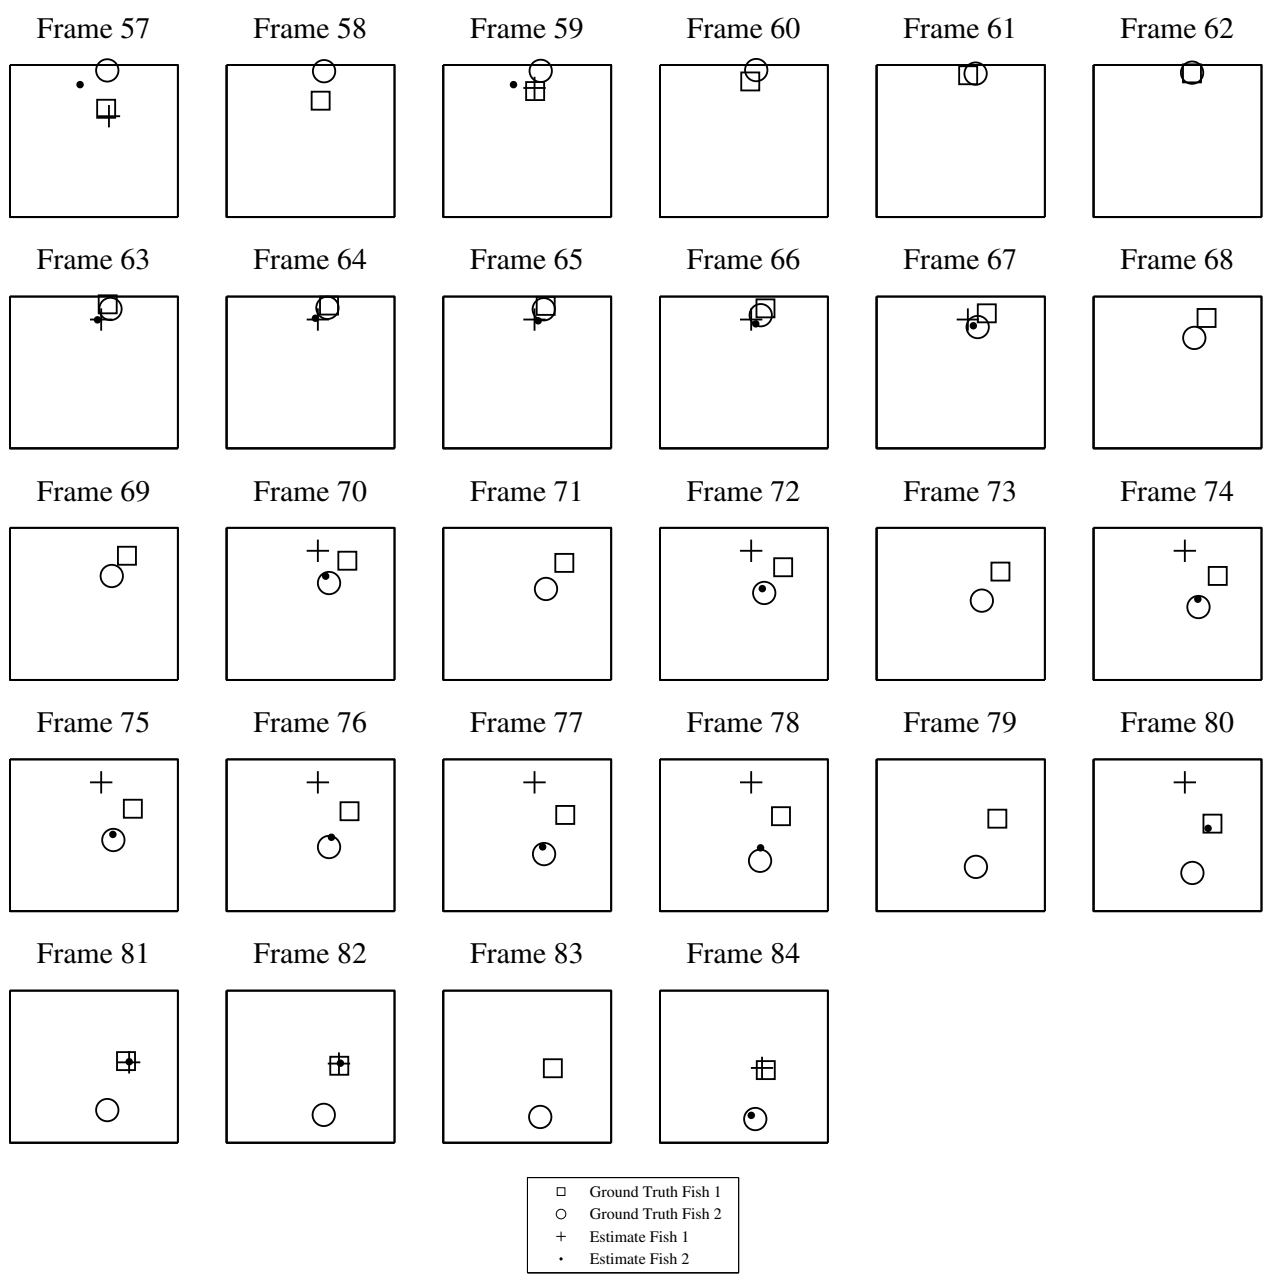


Figure 6. Complex interaction with change of direction and missed detections.