

QUT Digital Repository:
<http://eprints.qut.edu.au/>



Westall, Paul and Ford, Jason J. and O'Shea, Peter J. and Hrabar, Stefan (2008)
*Evaluation of Machine Vision Techniques for Aerial Search of Humans in
Maritime Environments*. In: Digital Image Computing: Techniques
and Applications (DICTA) 2008, 1-3 December 2008, Canberra, Australia.

© Copyright 2008 IEEE

Evaluation of Machine Vision Techniques for Aerial Search of Humans in Maritime Environments

Paul Westall^{*}, Jason J. Ford^{*}, Peter O'Shea^{*} and Stefan Hrabar[†]

^{*}*School of Engineering Systems, Queensland University of Technology, Brisbane, Australia
{p.westall, j2.ford, pj.oshea}@qut.edu.au*

[†]*Commonwealth Scientific and Industrial Research Organisation
ICT Centre, Brisbane, Australia
stefan.hrabar@csiro.au*

Abstract

Searching for humans lost in vast stretches of ocean has always been a difficult task. In this paper, a range of machine vision approaches are investigated as candidate tools to mitigate the risk of human fatigue and complacency after long hours performing these kind of search tasks. Our two-phased approach utilises point target detection followed by temporal tracking of these targets. Four different point target detection techniques and two tracking techniques are evaluated. We also evaluate the use of different colour spaces for target detection. This paper has a particular focus on Hidden Markov Model based tracking techniques, which seem best able to incorporate a priori knowledge about the maritime search problem, to improve detection performance.

1. Introduction

Human maritime search and rescue missions have always been a challenging task and an element of chance is involved in the detection of survivors [1]. Humans become fatigued and complacent after long hours of searching, reducing the chance of finding survivors. We propose the use of machine vision to automate the location of human survivors lost at sea, and present an evaluation of a number of techniques to achieve this.

Australia's search and rescue region alone covers approximately 53 million square kilometres, which is equivalent to nearly one tenth of the Earth's surface. This is a vast area to search and if robots, such as Unmanned Aerial Vehicles (UAVs) are employed, this could increase the probability of locating survivors. Given this search capability, robots would be able to assist and aid present manned search endeavours by

utilising UAVs as a force-multiplier. This would allow the current search efforts to be more flexible and to respond with appropriate force in a timelier manner.

A small number of studies into the automation of maritime searches have previously been conducted, yet all of these have restricted the search to small vessels and high visibility targets. Most notably, Sumimoto, et al [2, 3] have investigated the search for small bright orange vessels, such as life rafts, in the ocean. Methods such as morphological and high-pass filters were used to increase target signal-to-noise ratio (SNR) as well as contour matching to exploit shape information for distinguishing the target. However, only limited use was made of the colour information. In [4], Toet investigated the maritime search problem by choosing to fuse the morphological top-hat information from two different IR spectral frequency bands to reduce the affects of noise in the image produced by the surface of the ocean while searching for approaching kayaks.

Both of these investigations used images taken from static platforms looking out across the ocean, as from the bridge of a ship, searching for targets that are highly visible and largely above the surface of the water.

In light of this prior work, this paper will examine the performance of various point detection front-ends along with a number of track-before-detect techniques in an effort to find the most effective configuration for the detection of humans in an aerial maritime search environment.

This paper is organised in the following way: The next section contains a brief description of the problem and the constraints are outlined. Section 3 presents an overview of the proposed system, and Section 4 defines the colour transformations, Section 5 and Section 6 detail the point detection and tracking phase of the system, respectively. The experiments

performed are described in Section 7 and the results reported in Section 8.

2. Problem Definition

This paper specifically addresses the detection of a single person in the water during a daytime aerial search in a scenario where the survivor is without a high contrast floatation or location device. The images are taken from a downward looking camera and the ocean surface currents are no faster than 2m/s [5].

To maximise the effective search corridor of a single camera the proposed algorithm searches for persistent point-like targets. This feature of the proposed search approach also has a bearing on the flight altitude that the system can be flown at, given the resolution and field of view (FOV) of the camera.

Searching for a person in the water is difficult because the target may only occupy 1-3 pixels in a dynamic and constantly changing environment, and is therefore not easily distinguished from the background. Furthermore, the target has limited time within the field of view of the camera. Depending on the speed of the aircraft, the target maybe in view for as little as 3-5 seconds. Hence, a decision on the target's status must be made quickly to allow the operator to respond to the alarm.

3. System Overview

In point detection applications it is common to use techniques in the spatial domain such as spatial masking techniques [6] as opposed to methods such as Fourier transforms and the continuous wavelet transform that consider the entire image as a single entity.

Various temporal tracking techniques can then be used to discriminate the true target from noise using the target's properties such as size, shape, colour and temporal dynamics.

Therefore, the proposed detection system comprises two main components: a point target detection phase (front-end) to identify candidate pixels, followed by a temporal tracking phase that uses *a priori* knowledge and historical data to discern true target behaviour of the candidates (see Figure 1). We investigate four front-end image processing techniques and two temporal filtering approaches: Dynamic Programming (DP) and Hidden Markov Model (HMM) filtering.

4. Colour Spaces

Colour information can be useful for distinguishing foreground and background objects and when dealing

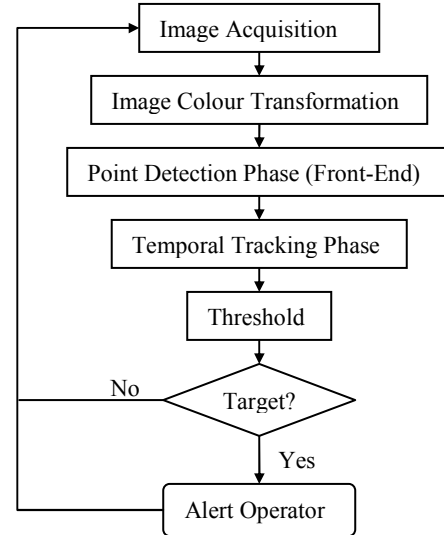


Figure 1. System Block Diagram

with small, pixel-sized targets all useful information should be exploited. Therefore, alternative colour spaces are investigated to determine which spaces maximise the SNR before other filtering stages.

The most likely part of the survivor's body to be seen above the water is their head. Hence, a provisional database of hair colours was compiled that were then modelled and combined to form target probability density functions (PDFs) to be used in the HMM. The dynamic programming approach does not allow for this kind of *a priori* knowledge to be directly incorporated into its' structure. PDFs for the ocean colour (i.e. background) were also established.

Based on the conventional Red-Green-Blue (RGB) colour space where the images were acquired, the target and background models lacked separation making it difficult for the HMM to consistently discern true targets. Therefore, alternative colour spaces were explored in an effort to isolate, as much as possible, target colour models from that of the background.

Using the inbuilt colour transformation functions in Matlab® the RGB images were transformed to equivalent Hue-Saturation-Value (HSV), Luma-Blue-Chrominance-Red-Chrominance (YCbCr) and Luma-In-Phase-Quadrature (YIQ) images, in an effort to exploit the useful properties of these alternate colour spaces.

Shadows often cause significant changes in intensity but have minimal effect on chromaticity. Some colour spaces, such as RGB, are more susceptible to illumination changes, while others, like HSV, can reduce the influence of the illumination changes, reflections and shadows [7].

5. Point Target Detection Phase

Small targets are difficult to detect because they are not easily distinguishable from noise and/or clutter. In the proposed system, the automated search would be conducted at an altitude that would result in an average human head [8] occupying only 1-3 pixels of the search camera's FOV, eliminating any shape information. For these reasons we investigate point detection techniques.

5.1. Mathematical Morphology

Morphology is based on using a structuring element to perform two fundamental operations, dilation and erosion, with combinations of these operations creating the open and close functions, described as follows:

$$\text{Opening} \quad A \circ B = (A \ominus B) \oplus B \quad (1)$$

$$\text{Closing} \quad A \bullet B = (A \oplus B) \ominus B \quad (2)$$

Morphological filtering of images, either Electro-Optic (EO) or Infra-Red (IR), has proven to be quite successful at discovering point-like objects in areas of aircraft collision avoidance [9] and multi-spectral IR target detection [4].

The particular morphological filtering implementation of the close-minus-open (CMO) technique presented in this paper is a consolidation of the filtering method employed by Casasent [10] and the filter application of Deshpande [11]. Four 1D slit-shaped structuring elements (one vertical, one horizontal, one on the leading diagonal and one on the trailing diagonal) are applied at both the close and open steps. This quad-filter approach allows only targets which are compact in all directions (i.e. point-like) to be enhanced. Non-compact clutter, such as white caps and large floating debris, are attenuated.

However, it was found that the basic CMO operations give rise to images that no longer correctly represent the zero mean nature of the image noise (which in turn degrades the effectiveness of the tracking phase) [9]. The following alteration to the CMO stage was made to preserve the sign of the CMO image output:

Using the preserved sign method of CMO reduces false alarms by an average of 20% as opposed to the normal CMO process [9].

$$\begin{aligned} CMO_{signed} &= (F_{in} - (F_{in} \bullet SE)) + (F_{in} - (F_{in} \circ SE)) \\ CMO_{signed} &= 2F_{in} - ((F_{in} \bullet SE) + (F_{in} \circ SE)) \end{aligned} \quad (3)$$

where CMO_{signed} = Signed Output Image

F_{in} = Input Image

SE = Structuring Element

5.2. Order Statistic Filters

Another widely used approach for finding small targets is to apply order statistic filters similar to those presented by Deshpande [11]. The order statistic filters, such as the median filter, are extensively used in multi-dimensional signal processing and are able to remove impulse noise, preserve geometrical features as well as being computationally efficient. For this study a mean-median filter was chosen and was implemented in a similar manner to that described in [11].

The mean-median filter is implemented in the following way:

$$y(m, n) = \text{mean}[z_1, z_2] \quad (4)$$

where $z_1 = \text{median}[x(m, n - N), \dots, x(m, n), \dots, x(m, n + N)]$

$z_2 = \text{median}[x(m - N, n), \dots, x(m, n), \dots, x(m + N, n)]$

and x = Input

y = Output of the Median window array of size $2N + 1$

where the focal pixel is (m, n)

Finally, the adaptive median filter (AMF) is an extension of the basic median filter (BMF), where the size of its window/kernel depends on the statistical properties of the neighbouring pixels of the focal pixel. Using this type of filter achieves two objectives: firstly, it thoroughly removes impulse noise and secondly, it reduces the extent of distortion that is evident with a comparable median filter.

As previously mentioned, these order statistic filters remove impulse noise; therefore the difference between the filtered image and the original produces an image of potential point-like targets which is then passed to the tracking phase.

6. Temporal Tracking Phase

Traditional temporal tracking algorithms often use thresholded measurements and attempt to associate track measurements over time – this method works well for high SNR scenarios, however becomes problematic once the target measurements are comparable to the noise and clutter. In a low SNR environment, lowering the threshold would degrade system performance by increasing false alarms.

Thresholding raw sensor measurements result in a loss of information. Keeping all measurement information allows for the potential to track low SNR targets in an unfavourable environment by integrating sensor responses over time to detect and track targets in high clutter and low SNR scenarios. However, clutter is reduced with the addition of the front-end, thus enhancing the ability to effectively identify plausible point-like targets.

Finally, a threshold is imposed as a final stage to the system to grant target status to those candidates that

have been tracked over a number of frames – this technique is commonly referred to as track-before-detect (TBD), and is a common concept in radar technology.

For the data considered in this study the search for targets between consecutive frames is restricted to within a 3x3 kernel as the target’s dynamics are limited within the maritime environment, as described below.

Assume the use of a 1024x768 pixel camera operating at 15 frames per second (fps) at a nominal search height of 150m above sea-level and travelling at 150km/h [1]. The highest documented surface current speeds (barring extreme environmental anomalies) are in the vicinity of 2m/s [5] which is approximately 7km/h. In addition to this surface current, the maximum swimming speed of a human is roughly 8km/h [12]. Aggregating these two velocities still does not exceed the maximum velocity allowed by our target dynamics model.

To see this, suppose the FOV of the camera in the direction of the aircraft x-axis is 60°, this makes each pixel equivalent to 0.226m (average human head diameter [8]). Thus, for a target to move at least two pixels, relative to the ground, between consecutive frames it would have to be travelling at a minimum speed of approximately 25km/h – a speed that is highly improbable for the type of targets of interest. As a result, for the problem of airborne maritime searches the slow-moving nature of the target allows the discrete velocity space to be limited to ± 1 pixel in both the x- and y- directions.

6.1. Dynamic Programming

Dynamic Programming is a widely used TBD method initially developed by Bellman in the 1940s and 1950s as a method of solving multistage decision problems. Gandi [13] and Carnie [9] applied DP to the image based target detection/tracking problem to optimise the target classification decisions at each frame (i.e. stage) of the process while using relatively simple constraints.

In this paper the target velocity is assumed constant, although small accelerations/manoeuvres are tolerable along plausible target trajectories. These target trajectories are modelled as state transitions and each transition is considered equally likely.

Assuming constant velocity, it can be shown [14] that for each discrete target state (i,j,u,v) at frame k , there are four possible state transitions corresponding to frame $k+1$, where (i,j) denotes a discrete position in the 2D image space and (u,v) denotes one of the four 2D velocity branches. Possible transitions are discretised into combinations of up/down and left/right pixel movement sectors. Consequently, only four

velocity branches are required to sufficiently describe the possible target motion, illustrated below in Figure 2.




Figure 2. DP state transitions

The dynamic programming algorithm used in this paper is divided into three stages – Initialisation, Recursion and Decision, and is based on the approach used by Gandhi.

Initialisation: The image created at the completion of each iteration, for each velocity branch (u,v) , is denoted by $F_{uv}(i,j,k)$. Initialisation of F_{uv} is:

$$F_{uv}(i,j,0) = 0 \quad \forall i,j,u,v \quad (5)$$

Recursion: The following expression is used to produce the intermediate image $F_{uv}(i,j,k)$.

$$F_{uv}(i,j,k) = (1 - \alpha)f(i,j,k) + \max_{u,v} (F_{uv}(i,j,k-1)) \quad \forall i,j,u,v \quad (6)$$

Where $f(i,j,k)$ is the image at frame k and the forgetting factor, $0 \leq \alpha \leq 1$, gives greater significance to past values as it is increased. Using the kernels defined by (u,v) , $\max_{u,v} (F_{uv}(i,j,k-1))$ is calculated with its focal pixel at (i,j) , see Figure 2.

By summing the frames the system is able to attenuate the affect of noise spikes that appear in the output of the CMO morphological filter within a few iterations. Therefore, only slow-moving persistent targets are maintained while transient anomalies, such as noise, are mitigated.

Decision: Immediately after the production of an intermediate image frame, $F_{uv}(i,j,k)$, from each velocity branch a single output frame is created using the maximum value of all branches on a pixel-by-pixel basis.

$$F_{\max}(i,j,k) = \max_{u,v} (F_{uv}(i,j,k)) \quad \forall i,j,u,v \quad (7)$$

Finally, a binary form of the $F_{\max}(i,j,k)$ image is established using a single threshold that is empirically chosen to avoid the tracking image noise floor, with pixels exceeding the threshold classified as targets.

6.2. Hidden Markov Model

The Hidden Markov Model is a powerful statistical tool involving stochastic processes that can be represented as an underlying discrete-value Markov chain state process that is partially observed through a sequence of measurements. HMMs have found use in many areas, such as signal processing, in particular speech recognition and document character recognition applications.

A HMM is characterised through the following:

- The set of Markov chain states $S = \{s_1, s_1, \dots, s_N\}$, where N is the total number of valid states for the model; and
- The HMM parameter set $\lambda = (\pi, A, B)$, where:

π is the initial state distribution vector (also known as *prior probabilities*), e.g. π is the probability of state i at the arbitrary time $t = 0$.

A is the state transition matrix (also known as *transition probabilities*), where $A = [a_{ij}]$, with a_{ij} being the probability of transition to state j given current state i .

B is the output distribution matrix (also known as *emission probabilities*), where $B = [b_{ik}]$, with b_{ik} being the probability or likelihood of observing feature k given current state i .

In our approach the number of states, N , used in the HMM is equal to the number of pixels of the input image (representing each of the possible locations of the target).

After being initialised (8), The HMM filter essentially acts as a recursive algorithm which evaluates the probability of a target being in each location given all previous observations (9). The output of this recursion algorithm is thresholded to determine if a target is present.

$$\text{Initialise} \quad \delta_1(i) = \pi_i \cdot b_{i,x_1}, \quad i = 1, \dots, N \quad (8)$$

where $N = \text{Number of states}$

$\pi_i = \text{Prior probability of being in state } s_i \text{ at time } t = 1$

$$\text{Recursion} \quad \delta_t(j) = \sum_{i=1}^N (\delta_{t-1}(i) \cdot a_{ij}) \cdot b_{j,x_t}, \quad (9)$$

where $2 \leq t \leq T$ & $1 \leq j \leq N$

$t = \text{current time step}$

$x_t = \text{observation at time } t$

The likelihood of a transition from one state, or pixel, to any of the surrounding eight states is defined by a discretised Gaussian surface with its peak located

on the current state and its standard deviation equal to one third of its operational extent.

In order to populate the emission probabilities matrix, B , the ratio of the target to background PDFs at each greyscale level was calculated. The target PDFs were generated by sampling a variety of naturally coloured hair specimens (i.e. black, honey blonde, brown, red and white hair) while the background PDF was assembled by sampling image frames of the test data that did not contain valid targets and approximating this data using a Gaussian curve.

Also incorporated into the emissions matrix is an augmented version of the information produced by the image processing front-end. The greyscale readings given by the front-end are mapped to a cumulative normal curve to enhance the point-like interpretation of the front-end filter as dim point-like targets can be somewhat disadvantaged in the typical filtering process. The parameters of the cumulative Gaussian curve are regulated by the statistical properties of and separation between the target and background PDFs.

These two metrics, the ratios of hair colour PDFs to background PDF and the reinterpreted front-end readings, are then combined to form a single emissions matrix to be used in the HMM.

7. Experiments

Below we describe trials that were conducted on both simulation and real flight data to gauge the performance of various system configurations, and to determine the most effective approach.

7.1. Simulations

A simulated ocean scene was created using Blender, an open-source 3D animation application. The scene was animated to generate image sequences similar to those that would be produced during an aerial search operation, and these were used to test the detection algorithms.




Figure 3. Blender image samples

Blender allowed us to vary wave height and speed, the colour of the ocean, the amount of white water as well as the reflections created by the sun. Two

scenarios were used – calm water and wavy white water (left and right images in Figure 3, respectively).

The scene was made to scale and the images shown in Figure 3 were created from a simulated camera height of 1000ft. A human figure was also inserted into the scene with hair colour matching the data used to populate the emissions matrix in the HMM.

Ten data sequences were generated from the two different sea states and five hair colours.

7.2. Flight Data

A flight test was conducted in a Cessna 172 over the beaches of the Gold Coast, Australia during June 2007. Images were captured from an altitude of approximately 500ft at 80 knots by a downwards-pointing camera mounted to the wing strut. A Point Grey Flea[®] camera fitted with a 185° FOV Fujinon YV2.2X1.4A-2 fisheye lens produced 1024x768 images at 15Hz. Camera pose (based on GPS and IMU) were also logged for each frame

Although data was captured with a fish-eye lens, the region in the centre of the image has high spatial resolution without much distortion compared to the areas towards the periphery of the image. The centre portion of the original 1024x768 image was cropped to produce a 267x200 image. No further image rectification was performed.

A sample image frame from the collected data is shown in Figure 4. A target is visible near the top-left corner (a surfer in a red rash shirt/wetsuit). This target remains within the cropped 267x200 image sequence for approximately 2-3 seconds, providing a suitable data set.




Figure 4. Flight data image sample with target

Note that for both the simulated and real data sequences, the images were post-processed to compensate for camera motion before being used by the detection algorithm.

7.3. Performance Metrics

A set of metrics were used to compare the performance levels of the system. The performance of the system was quantified using the following metrics:

FAR – False Alarm Rate is the average number of false alarms per image.

MDR – Missed Detection Rate is the average number of missed detections of true targets per image sequence.

FFTT – First Frame of True Target Detection (true target first appears in frame 22 in simulated data sets).

FA length – False Alarm track length is the average number of consecutive frames that false alarms are considered targets.

Signal-to-Noise ratio of the image is calculated according to the following:

$$SNR = 10 \times \log_{10} \left(\frac{\text{mean}(\text{signal pixel intensity}^2)}{\text{mean}(\text{noise pixel intensity}^2)} \right) \quad (10)$$

8. Results

Simulation-based tests were performed with 96 different system configurations (four front-ends, two TBD and four colour spaces – three layers each) and 255 threshold levels. Representative subsets of the results are shown in Figures 5-7 and Tables 1-3 below. The best performing configurations were then tested on the real data (shown in Figure 8).

8.1. Sea State

Figure 5 illustrates that white waves on the surface of the water reduce the detectability of the target by introducing clutter that is of the same colour as potential targets. On the peripheries of some of the wave crests the white water diffuses to form clutter similar to the size of the target, complicating the search.




Figure 5. Detection under different sea states

The graph in Figure 6 shows that white colour targets are more difficult to find as they are mistakenly grouped with the white water clutter by the front-end.




Figure 6. Detection of different target colours

8.2. System Component Evaluation

The effect of the different colour spaces on the performance of the system is measured by the SNR after the point detection front-end, outlined in Table 1. The HSV colour space is the clear choice with more than 10dB greater performance than alternative colour spaces.

The SNR of the front-end (SNR-PNT) is used to evaluate to what extent the front-end can reduce clutter and enhance the target. Table 2 reveals that the most effective point detection method assessed is CMO, with a strong improvement over other techniques.

Table 1. Top performing colour space layers

Colour Layer	SNR-CS (dB)
HSV1	39.0149
YIQ2	27.8894
HSV2	27.4568
YCC3	26.8417
RGB1	26.5925

Table 2. Front-end SNR

Front-End	SNR-PNT (dB)
CMO	26.6827
AMF	24.1242
MnM	22.1189
BMF	19.6493

Table 3. Tracking phase performance

Tracker	SNR-TRK	FA Length	FFTT
DP	20.5936	2.5704	24.6594
HMM	31.6800	1.8255	26.7432

The tracking phase SNR (SNR-TRK) reveals that the HMM is far more effective at distinguishing targets in this environment. This is also supported by HMMs lower FA length. However, DP is able to detect and classify the true target approximately 2 frames (0.13 seconds) earlier than HMM, see Table 3.

8.4. Overall System Configuration




Figure 7. System configuration performance

There are 96 system configurations; some of the top performers are shown in Figure 7. The first term of the curve identifier is the colour space and layer used, the second term is the front-end and the third term is the tracking method. These results agree with the performance assessment of each of the individual phases, however the effectiveness of BMF is unexpected based on the results in Table 2.

These five configurations were then applied to the real data to produce the subsequent graph, Figure 8.




Figure 8. Detection performance of real data

This shows that the HSV1-CMO-HMM configuration produced the best performance, supporting the findings described in Table 1 and Table 2. Note however, that the background PDF used for the real data was different to that of the simulated data, as the colour distribution of the ocean between the sequences are not identical.

These tests have shown that the colour distribution input of the ocean sequence has a significant bearing on the performance of the system. As there is no single configuration that consistently outperforms the rest, we propose that perhaps fusing the output of a variety of colour transformations may make more effective use of the information, improving overall system performance.

9. Conclusions

We have evaluated and compared a number of vision-based techniques to aid in the detection of human survivors in maritime environments. Our investigation included evaluating four different colour spaces, four different point target detection techniques, and two different temporal tracking techniques. All possible combinations of the above were tested on synthetic image sequences.

This evaluation showed that the colour space layers most suited to the maritime scenario are HSV1, YCC3 and YIQ2. Of the point target detection techniques tested, CMO and BMF were shown to be most effective at reducing the effect of ocean clutter and improving target signal. Additionally, our study suggested that Hidden Markov Model based temporal feature tracking outperformed Dynamic Programming as it was best able to exploit *a priori* knowledge of the environment. We also propose that combining the output of a number of image colour transformations may produce the best overall system performance.

Although the system performance of the currently proposed approach is not suitable for real search operations (approximately one false alarm per frame for an acceptable MDR), this study was successful in assessing the preliminary performance of different system configurations. The best performing of these worthy of more thorough investigation.

The results illustrate that machine vision is a viable technology in the maritime human search application. It has the potential to play an important role in manned aircraft searches by drawing attention to possible areas of interest.

10. Acknowledgements

Flight data in this study was obtained courtesy of the Australian Research Centre of Aerospace Automation (ARCAA). Computational resources and services used in this work were provided by the High Performance Computing and Research Support Unit, Queensland University of Technology, Brisbane, Australia.

11. References

- [1] Australian National Search and Rescue Council, "National Search and Rescue Manual," Version 0501 ed Australia: Australian Maritime Safety Authority, 2005.
- [2] T. Sumimoto, K. Kuramoto, S. Okada, H. Miyauchi, M. Imade, H. Yamamoto, and T. Kunishi, "Machine Vision for Detection of the Rescue Target in the Marine Casualty," in *Industrial Electronics, Control and Instrumentation, 1994. IECON '94., 20th International Conference on*, 1994, pp. 723-726 vol.2.
- [3] T. Sumimoto, K. Kuramoto, S. Okada, H. Miyauchi, M. Imade, H. Yamamoto, and Y. Arvelyna, "Image Processing Technique for Detection of a Particular Object from Motion Images," in *Industrial Electronics, 2001. Proceedings. ISIE 2001. IEEE International Symposium on*, 2001, pp. 1662-1666 vol.3.
- [4] A. Toet, "Detection of Dim Point Targets in Cluttered Maritime Backgrounds through Multisensor Image Fusion," in *Targets and backgrounds. Conference No8*, Orlando, FL, 2002, pp. 118-129.
- [5] K. Ouchi, "The effect of SAR bandwidth ratio and current variation on ocean current measurements by along-track SAR interferometer," in *Geoscience and Remote Sensing Symposium, 1994. IGARSS '94. Surface and Atmospheric Remote Sensing: Technologies, Data Analysis and Interpretation., International*, 1994, pp. 723-725 vol.2.
- [6] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, Second Edition ed. New Jersey: Prentice Hall, 2002.
- [7] Y. Ren and C. S. Chua, "Bilateral learning for color-based tracking," *Image and Vision Computing*, vol. In Press, Corrected Proof, 2008.
- [8] T. Partala, V. Surakka, and J. Lahti, "Affective Effects of Agent Proximity in Conversational Systems," in *Third Nordic Conference on Human-Computer Interaction*, Tampere, Finland, 2004, pp. 353-356.
- [9] R. Carnie, R. Walker, and P. Corke, "Image Processing Algorithms for UAV "See and Avoid"," in *2006 IEEE International Conference on Robotics and Automation*, Orlando, FL, 2006, pp. 2848-2853.
- [10] D. Casasent and A. Ye, "Detection Filters and Algorithm Fusion for ATR," *Image Processing, IEEE Transactions on*, vol. 6, pp. 114-125, 1997.
- [11] S. D. Deshpande, M. H. Er, V. Ronda, and P. Chan, "Max-Mean and Max-Median Filters for Detection of Small Targets," in *Signal and Data Processing of Small Targets 1999*, Denver, Colorado, 1999, pp. 74-83.
- [12] Infostrada Sports, "Swimming - World Record Progression Men 50m Freestyle," International Olympic Committee 8 July 2004.
- [13] T. Gandhi, M.-T. Yang, R. Kasturi, O. I. Camps, L. D. Coraor, and J. McCandless, "Performance Characterization of the Dynamic Programming Obstacle Detection Algorithm," *Image Processing, IEEE Transactions on*, vol. 15, pp. 1202-1214, 2006.
- [14] S. M. Tonissen and R. J. Evans, "Target Tracking using Dynamic Programming: Algorithm and Performance," in *Decision and Control, 1995., Proceedings of the 34th IEEE Conference on*, 1995, pp. 2741-2746 vol.3.