Image Processing Techniques for Metallic Object Detection with Millimetre-wave Images

C. D. Haworth *, Y. R. Petillot and E. Trucco

Electrical, Electronic and Computing Engineering
School of Engineering and Physical Sciences
Heriot-Watt University
Edinburgh, EH14 4AS
United Kingdom

Abstract

In this paper we present a system for the automatic detection and tracking of metallic objects concealed on moving people in sequences of millimetre-wave images, which can penetrate clothing, plastics and fabrics. The subjects are required to enter one at a time and turn round slowly to ensure complete coverage for the scan. The system employs two distinct stages: detection and tracking. In this paper a single detector, for metallic objects, is presented which utilises a statistical model also developed in this paper. Target tracking is performed using a Particle Filter. Results are presented on real millimetre-wave image test sequences and indicate an excellent rate of success for threat identification. Encouraging results for target tracking are also reported.

Key words: millimetre-wave, detection, tracking, crime prevention

1 Introduction

We present a system for the automatic detection and tracking of metallic objects concealed on moving people in sequences of millimetre-wave (MMW) images, which can penetrate clothing, plastics and fabrics. The subjects are required to enter one at a time and turn round slowly to ensure complete

* *E-mail*: C.D.Haworth@hw.ac.uk
Telephone: +(44) (0)131 451 3299
Fax: +(44) (0)131 451 4155

coverage for the scan. The long-term aim of this work is to develop a MMW system to enhance the public security screening process for concealed threats. For example, airport passenger screening which routinely employs metal detectors. This papers focuses on the detection of concealed metallic objects, presenting the performance in both detection and tracking. At this stage no attempt to implement a real-time system has been made but because real-time operation would be a requirement in the final application the computational complexity of the algorithms developed is taken into account.

MMW imaging is emerging as an important modality for security and surveillance thanks to recent advancements in MMW sensing technology. Providing full monochrome images highlighting concealed threats opens the possibility to analyse shape and locate threats on the body, which is far beyond the reach of conventional metal detection portals. A recently demonstrated proof-of-concept sensor developed by QinetiQ (Sinclair et al., 2001) provides video-frame sequences with near-CIF resolution (320×240 pixels) and can image through clothing, plastics and fabrics. The combination of image data and through-clothes imaging offers huge potential for automatic covert detection of weapons concealed on human bodies via image processing techniques. Previous trials of the QinetiQ MMW sensor, involving the Department of Transport and British Airport Authority (BAA), showed potential for passenger screening at airports (Murphy et al., 2002), public event security (Sinclair et al., 2001) and detection of illegal passengers in lorries. All trials involved human operators.

The sequences in this paper are generated by an electro-optic sensor working between infra-red (IR) and microwave wavelengths. The sensor forms an image of the temperature received from the scene, which is a standing human subject turning around slowly. Figure 1 shows examples of frames from a typical sequence considered in our work. A person turns around by 360° in front of the sensor and is captured at video rate (12 frames per second). The person is allowed to turn at their own speed, to allow for various levels of physical ability but this normally takes approximately ten seconds (120 frames). The temperature (and therefore the pixel intensity) is a function of the reflectivity, emissivity and transmissivity of the scene surfaces. At the wavelength used, metallic objects tend to appear bright as they are highly reflective, the human body less bright as it is partially reflective, and clothes partially transparent. An illumination chamber is required for indoor operation (Coward and Appleby, 2003) but does not expose the subject to harmful radiations.

To the best of our knowledge, very little work has been reported on the automatic analysis of MMW sequences or images. Basic image segmentation (Slamani et al., 1999; Keller et al., 2000) has been reported with some success. Shape identification on the segmented images (Slamani and Ferris Jr., 2001) has been investigated and suitable shape descriptors proposed. However,

Fig. 1. Example MMW sequence showing a human subject. Notice the speckle noise pattern particularly apparent on the torso. Substantial smoothing which is applied during image formation to minimise visual artifacts.

the image quality is poor with a small field-of-view and cannot be gathered at frame-rate. Due to these limitations only a very constrained stationary scene is considered. The proposed shape descriptors prove reliable under the operational constraints. In comparison, the QinetiQ sensor used in this paper is a real-time, head-to-toe sensor with considerably improved image quality. More recently, some work on the detection and segmentation of metallic objects has been proposed (Haworth et al., 2004, 2005) for the QinetiQ sensor. Work has also been presented on modelling image formation in the QinetiQ MMW sensor (Grafulla-Gonzalez et al., 2005). Alternatively, image fusion (Varshney et al., 1999; Xue and Blum, 2003) has been employed on MMW / Infra-Red. While this work has produced good quality images for human operators, no work has been reported on automatic detection in fused MMW-IR images and it is not clear that there would be any benefit. The main contribution of our work is therefore to apply advanced image processing techniques for the automatic detection of concealed weapons to a new video imaging modality of potentially high value for many applications in public security.

This paper is organised as follows: Section 2 presents a statistical mixture model. The mixture model is then used in Section 3 to develop a classification strategy for millimetre wave images. The classification first examines the sequence, then individual frames and finally regions of a frame to identify possible threats. In Section 4 target tracking is presented that is designed to improve the robustness of the threat detection and provide extended information on the position of the threat on the subject. Results are presented in Section 5 on real MMW sequences.

2 Mixture Models for MMW Images

MMW images offer good data for material discrimination as different materials yield, generally speaking, different image properties. In analysing the image

statistics it would be desirable to have an understanding of the full physical process which could be incorporated in a model for the MMW image formation process. However, given the complexity of the MMW imager and the extensive amount of hardware calibration, software equalisation and interpolation undertaken to produce a MMW image, this is an unfeasible task.

In this paper we adopt an approach modelling the differences in image properties statistically, using a weighted mixture model in which each pdf, f_i, is associated to a specific material:

$$f_{mix} = \sum_{i=1}^{N} \alpha_i f_i(\boldsymbol{\theta}) \tag{1}$$

where α_i is a weight and $\boldsymbol{\theta}$ a vector of parameters.

To identify the best-fitting pdf for each material (incl. background, i.e., non-figure pixels), we built a number of mixture models made by combinations of standard distributions (e.g. Gaussian, Rayleigh, Laplacian), optimised the parameters with a standard Maximum Likelihood (ML) algorithm and picked the best fitting combination for the observed image histograms using a Chi-Square test. We started with background-only sequences (no subject) to identify the background distribution. We then moved to sequences of scenes with a subject but no threats, then with a subject carrying threats (metallic objects). The final result is a best-fitting mixture model for each material (types of component distributions and parameters).

As a qualitative example, Figure 2 shows four histograms. The results of the ML distribution fit for a scene containing a subject carrying no threats are shown on the top row. Here, a two-component mixture model is used: two Gaussians, leading to poor fit, and Laplacian-Rayleigh, showing good fit and little overlap between component distributions. The results for a subject scanned carrying a metallic object are shown on the bottom row. In this case it is clear that a two class model (bottom left) is insufficient and that the model must be adapted to continue to provide a good fit. The inclusion of a specific PDF for metal improves the fit (bottom right). In particular it significantly reduces the overlap between the background and body PDFs. For these reasons, the identification of frames containing metallic objects is vital to allow the correct mixture model to be employed.

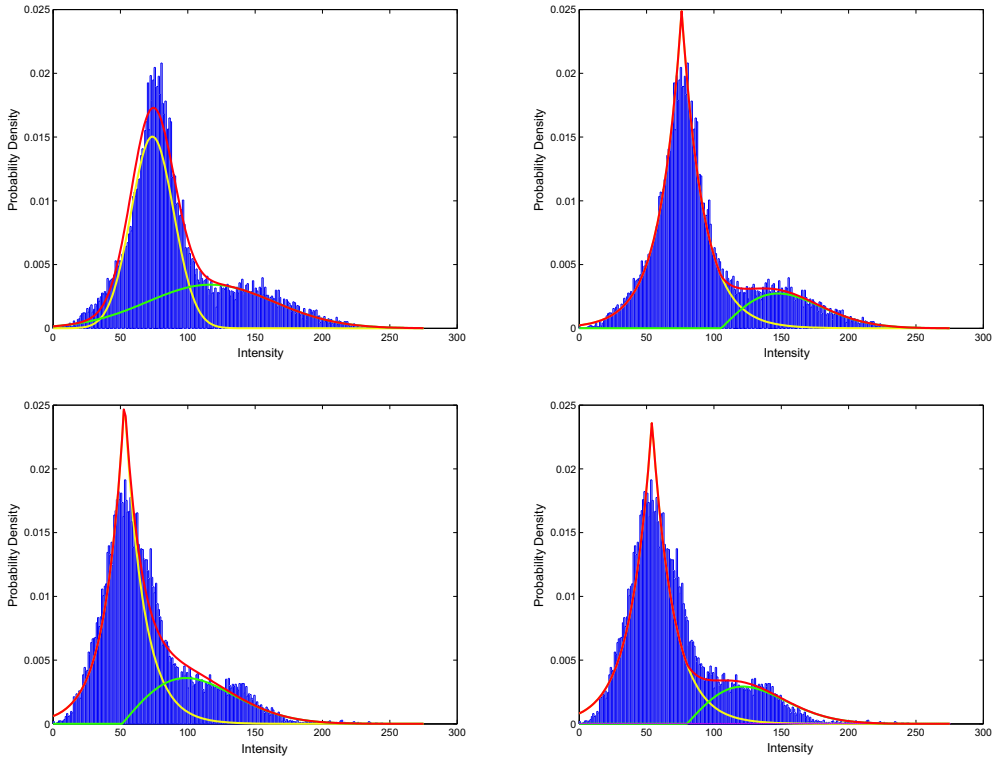

Fig. 2. Example of PDF fitting. A subject screened with no metallic objects fits well with a Laplacian and Rayleigh (top right) but poorly with two Gaussians (top left). However, a Laplacian and Rayleigh fit is not correct for a subject scanned with a metallic object (bottom left) and requires the addition of a PDF to account for the metal (bottom right). Note that the metal PDF is Uniform with a low probability making it difficult to see, however, its effect can be seen on the overall fit of the mixture model.

3 Classification of MMW Images

The development of the classification scheme is broken down into three stages. The identification of sequences containing threats is addressed first (Subsection 3.1 and then the identification of individual frames containing threats (Subsection 3.2). Finally an approach for the identification of specific threat regions within each frame is presented (Subsection 3.3).

3.1 Identifying sequences containing threats

The presence of metallic objects changes the maximum temperature recorded significantly, providing a good criterion to identify frames containing threats. Within a sequence, the range of variation of the maximum image temperature provides a reliable measure of the presence of a threat when compared

to a normalised threshold. After careful analysis of the test data set it was determined that a single threshold, set on the range of variation of maximum image temperatures, was suitable:

$$\frac{MAX(I_t^{max}) - MIN(I_t^{max})}{MAX(I_t^{max})} > 0.03 \tag{2}$$

where I_t^{max} is the maximum intensity (temperature) for frame t. In contrast, detecting which frames in the sequence contain objects is more difficult and this is discussed in the next section. In theory, with the subsequent method for identifying frames (Subsection 3.2) it would not be necessary to use this approach to identify if the sequence contains a threat. In practice, however, this technique is a computationally cost-effective method of determining if each frame needs to examined individually for metallic objects.

3.2 Identifying frames containing threats

To solve the problem of identifying individual frames containing metallic objects we trained a Hidden Markov Model (HMM) (Rabiner, 1989) to detect significant changes in maximum temperatures (i.e., image intensities). The system was trained on 3 different sequences (a tall man, a female, a short man). In our implementation, the data is first quantised into 10 levels and the hidden field is composed of 2 states (threat,no threat). A Baum-Welch algorithm (Rabiner, 1989) is used for parameter estimation, and a Viterbi algorithm (Rabiner, 1989) to determine the optimal state sequence. Figure 3 shows the maximum temperature signal for six sequences, with the variation in level, range and type of signal obvious.

While it would have been possible to approach this problem using empirically determined rules and thresholds, as demonstrated in Section 3.1, this would have produced a complex decision that would have required continual updating. The HMM approach avoids these explicit rules and thresholds. Furthermore, the HMM could easily be retrained on extended data sets should a specific problem present itself. Retraining might be justified if changes in the sensor's environment was significantly changing the thermal response of the sensor. However, this has not proven necessary so far.

3.3 Locating threat regions within frames

We now turn to the problem of locating the image region corresponding to a metallic object in frames classified as containing threats. It is important to

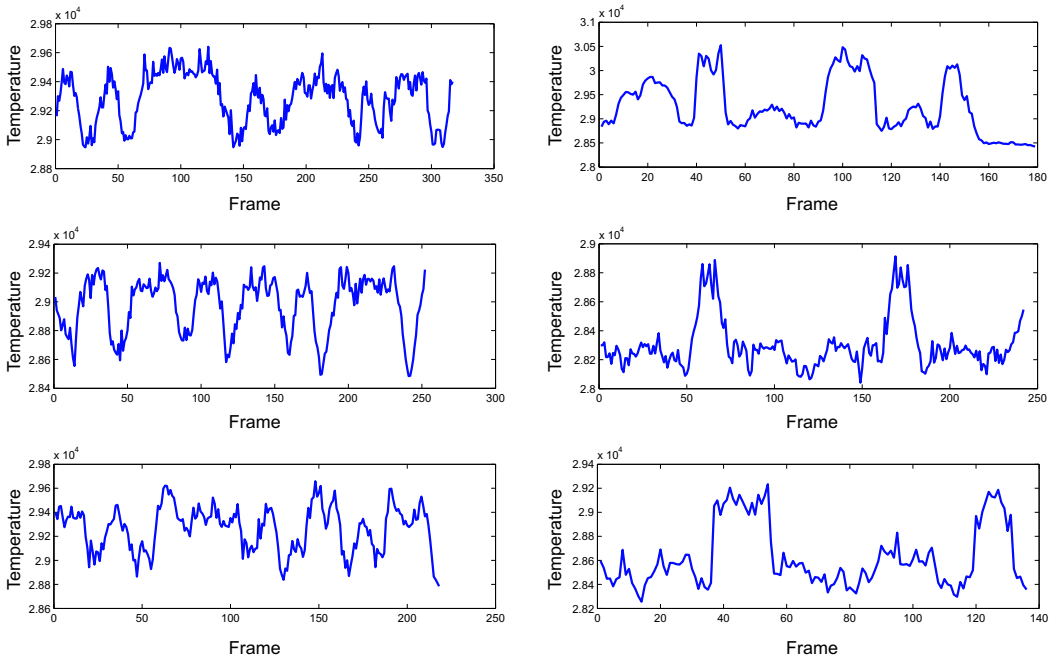

Fig. 3. Examples of the variation in level, range and type of signal found in different MMW sequences. The three sequences on the left do not have a metallic object present. The three sequences on the right do have a metallic object. All frames in the six examples were correctly classified by the HMM.

remember that the target application could require close to real-time operation and that techniques relying on extensive computing resources may not be suitable. For this reason, two alternative techniques were considered: k-Means and Expectation Maximisation (EM). The k-Means clustering technique can be viewed as a simplified variant of the EM clustering technique. This provides us with a direct trade-off between accuracy and efficiency.

The first technique, k-Means unsupervised clustering, uses a simple statistical measure, the sum of the squared Euclidean distances from the mean of each cluster, and has been well investigated in the literature for a range of applications (Kanungo et al., 2002). It has the distinct advantage of having low computational costs. However, k-Means is biased towards finding symmetric clusters and this cannot be guaranteed for the PDF mixture models selected in Section 2.

The second technique, Expectation Maximisation (Bilmes, 1997), uses a Maximum Likelihood Estimate to recompute the pdf parameters until a convergence criterion is met. We initialise the mixture model to the one containing the optimal distributions for the background-body-metal case (as defined in Section 2) with default parameters. Notice that this is not strictly necessary for the EM algorithm, but improves the convergence speed significantly.

An example of comparative results is shown in Figure 4. The results show a

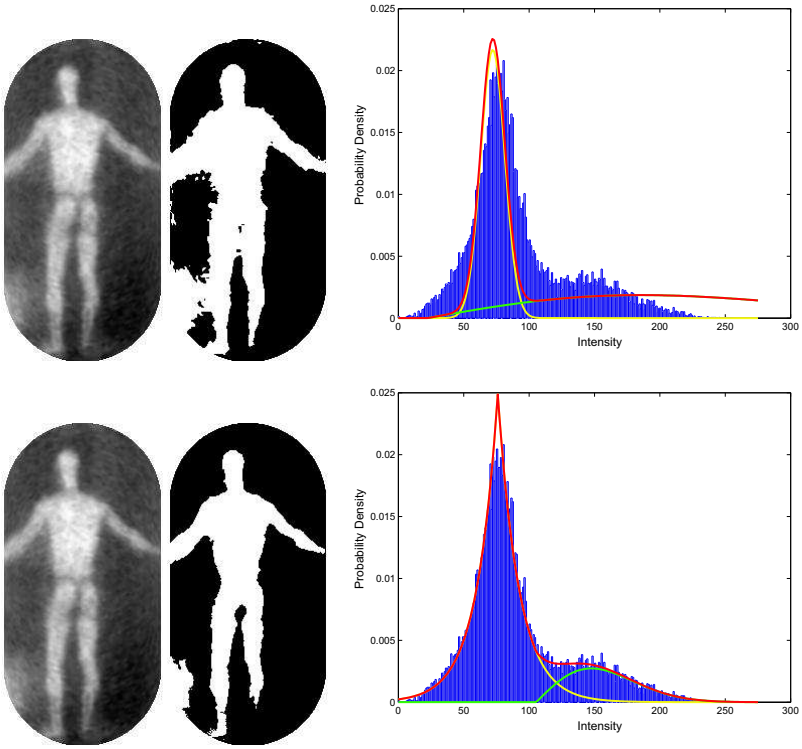

Fig. 4. Comparison between a k-Means (top) and EM (bottom) clustering, showing the original image, resulting image and the PDF overlaid onto the histogram.

typical case from our experiments, with k-Means providing a poorer PDF fit in comparison to EM. The resulting image is an acceptable approximation to the EM clustering with some degradation in performance. This reflects the theoretical trade-off between accuracy and performance, as EM is considerably more expensive than k-Means. Further work on approximate techniques, such as the k-Means, may lead to an optimal trade-off should computational complexity become a significant issue for system implementation. However, extensive field trials would be required to establish if acceptable levels of performance are met.

An example of working threat location is shown in Figure 5 using EM clustering, where the estimated threat region is highlighted in white.

4 Tracking threat regions

4.1 Problem Introduction

The results of the classification stage applied to sequences of persons carrying metallic threats is twofold: a set of frames showing possible metal threats, plus,

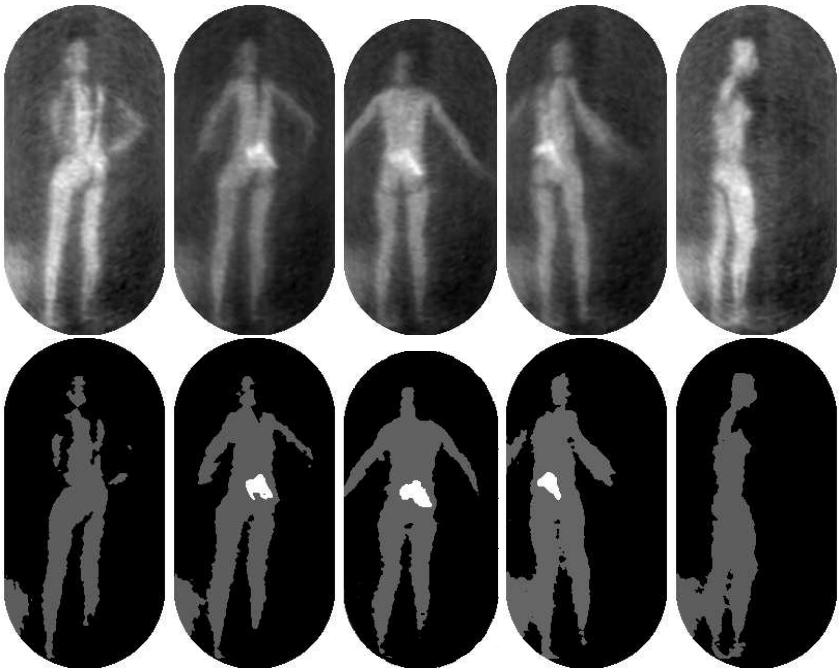

Fig. 5. Expectation-Maximisation segmentation (bottom row) of a scene containing a potential target in a MMW image (top row). The frame identification has worked well with the correct number of classes employed in the classification at all stages.

in each such frame, the regions corresponding to possible threats. Such regions are characterised by frame number, centroid, and area. To obtain further robustness and to provide a reliable estimate of the location of each threat on the subject, the threat regions can be now be tracked throughout a sequence for as long as the region remains visible. The subject is requested to undertake a constrained motion and, while we do not know their true motion, it is this constraint on the subject's motion combined with the tracking that can help us to achieve the accurate and reliable performance necessary for a public security application. The problem is made more difficult by the noisy nature of the MMW images being operated on.

Tracking objects in visible-wavelength sequences is a well-studied problem in image processing and computer vision (Trucco and Plakas, 2005; Comaniciu and Meer, 2002; Comaniciu et al., 2003). Particle filters (PF) (Arulampalam et al., 2001) are a powerful class of algorithms removing the Gaussian-distribution constraint typical of Kalman filters. They also provide robustness against clutter, a significant problem in MMW images given the noise characteristics. A common problem with PF is the *degeneracy problem* where after several iterations all but a few particles have negligible weights. For this reason the Regularised PF (RPF) (Arulampalam et al., 2001) has been developed to have an improved re-sampling stage, helping to avoid the degeneracy problem by calculating the effective number of particles in the system and when necessary replacing particles with a small weight (ineffective particles) with a new

set of particles while maintaining a good distribution of particles across the state space.

4.2 Particle Filter Models

Target tracking with the RPF requires the definition of two models: a *system model* defining the evolution of the state with time and a *measurement model* which relates the measurements to the state. Associated with these models are noise sources: process noise v_t and measurement noise n_t. A detailed derivation of the inference problem for the PF can be found in (Arulampalam et al., 2001).

Using the current classification scheme (Section 3) the state vector can contain any of the following: centroid (position) ($\mathbf{x}$), velocity ($\dot{\mathbf{x}}$), area (ϕ) and average gray-level (intensity) (I). Through manual observation it was determined that the average gray-level provided limited consistent information. In contrast, the area component can provide important information to supplement the position and velocity components in providing a robust estimate. For these reasons, the state vector was selected to contain the position, velocity and area of the target:

$$x_t = \left(\mathbf{x_t} \ \dot{\mathbf{x}}_\mathbf{t} \ \mathbf{y_t} \ \dot{\mathbf{y}}_\mathbf{t} \ \phi_t \right)^T \tag{3}$$

For our application a linear Gaussian system model using the following standard state space model for a constant velocity model and random walk model for the area (see (Bar-Shalom and Fortmann, 1988), for example) proved suitable:

$$x_t = \begin{pmatrix} 1 & T & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & T & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} x_{t-1} + \begin{pmatrix} T^2/2 & 0 & 0 \\ T & 0 & 0 \\ 0 & T^2/2 & 0 \\ 0 & T & 0 \\ 0 & 0 & T \end{pmatrix} v_{t-1}, \tag{4}$$

and observation model:

$$o_t = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} x_t + n_t. \tag{5}$$

v_t and n_t are the process and measurement noise components respectively, which are uncorrelated, and o_t is the observation vector. Noise components are

assumed Gaussian as this proved sufficient for this tracking scenario. Note that this is not a restriction of the Particle Filter and alternative noise distributions could be employed in the future if required.

4.3 Particle Filter Tuning

All parameters were selected through detailed testing to provide optimal performance across all test sequences. For the number of particles used in the system a value of 2000 was determined to be the preferred trade-off between accuracy and computational cost. As a starting point for position and velocity the motion was observed to be slow and uniform (<3 pixels / frame) in the test sequences, suggesting corresponding prediction variances of 10. Due to the high level of noise within the image classification a higher level was required for the measurement variances, with a final value of 50 used. Finally, it is important to allow for the unpredictable area changes associated with our illumination environment, forcing the related covariance components to be significantly greater than those of position and velocity.

Using the state vector shown in Equation 3, the covariance matrices C_{vt} and C_{nt}, of process noise v_t and measurement noise n_t, employed in our experiments were:

$$C_{vt} = \begin{pmatrix} 10 & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 100 \end{pmatrix} \tag{6}$$

$$C_{nt} = \begin{pmatrix} 50 & 0 & 0 \\ 0 & 50 & 0 \\ 0 & 0 & 1000 \end{pmatrix} \tag{7}$$

5 Experimental results

5.1 Target Detection and Tracking Accuracy

To evaluate our system, eight test sequences were employed, four with subjects without a threat and four with subjects carrying a threat, giving a total of 1629 frames and including 137 frames where a threat is visible. Table 1 summarises

Table 1
Test Sequences Employed

Sequence	Frames	Threat	No. Threat Frames
Plain01	211	No	—
Plain02	252	No	—
Plain03	218	No	—
Plain04	236	No	—
Threat01	242	Yes	24
Threat02	155	Yes	27
Threat03	179	Yes	56
Threat04	136	Yes	30
Total	1629	4/8	137

the details of the test sequences. In each test sequence the subject is required to enter the sensor, stand with legs slightly apart and arms slightly raised before turning through 360 degrees. The crucial aspect is to avoid the subject occluding a specific body part from the sensor. The number of frames varies between sequences because subjects where allowed to rotate at their own speed and were allowed to include some minor unrequested motion, such as pausing or arm waving. A small handgun was used as the metallic threat and this was concealed under the subject's clothes in various positions. The number of threat frames is the total number of frames across the entire sequence that the threat is visible. The number of consecutive frames a threat is visible varies depending on placement. For example, threats on the front or back of the body are only occluded by the trunk but threats on the side can be occluded by both trunk and arms resulting in more frequent but short-lived occlusion.

Table 2 shows the results of the sequence threshold (Subsection 3.1) in column 2 and the HMM frame identification (Subsection 3.2) in columns 3-5. The HMM results give percentage error in classified frames ($Error$) with a breakdown of target frames missed (E_{miss}) compared to false alarms (E_{false}). The results clearly show that the two-stage threat identification performed very effectively. The missed target frames were primarily in situations where the target was identified through shape rather than intensity. Figure 6 provides an example. This suggests that a shape-based detector combined into the system could improve reliability and robustness of the system.

Finally Table 3 shows results for the EM classification and RPF target tracking, giving the average number of targets (true target + clutter) per frame for the sequence and RMSE of the tracked position. The ground truth for the target position was manually tracked and is accurate to ± 2 pixels. It can clearly

Fig. 6. Two example frames showing where the metallic threat object can be identified solely from the intensity (left) and where only the shape of the object identifies the threat (right - threat is on the subjects right hip).

Table 2
Threat Identification with the sequence threshold (Subsection 3.1) in column 2 and the HMM results (Subsection 3.2) in columns 3-5. The HMM results provide the overall error ($Error$) and the breakdown in errors due to false alarms (E_{false}) and errors due to missed frames (E_{miss}). Frames with a threat and false alarms (clutter) were considered to be correctly identified. The sequences identified as having no threat were no examined on a frame-by-frame basis.

Sequence	Threat?	$Error$	E_{false}	E_{miss}
Plain01	No	—	—	—
Plain02	No	—	—	—
Plain03	No	—	—	—
Plain04	No	—	—	—
Threat01	Yes	8%	0%	100%
Threat02	Yes	3%	0%	100%
Threat03	Yes	5%	22%	78%
Threat04	Yes	8%	0%	100%

be seen that excellent target tracking results have been achieved, even in the sequences with considerable clutter (Threat01, Threat02). The comparatively poorer tracking results seen in Threat02 are due to the very short time span over which the threat is visible (approx. 9 frames on each occasion compared to an average of 15 frames for other sequences). In this instance, the particle filter does not have enough time to converge. Some example frames from the RPF target tracking are shown in Figure 7.

Table 3
Target Tracking showing the average number of targets (true threats + false alarms) and tracking accuracy (Root Mean Square Error in pixels) for the true threat.

Sequence	Average Targets	RMSE
Threat01	2.4	8.1
Threat02	2.1	11.6
Threat03	1.3	5.1
Threat04	1.1	5.5

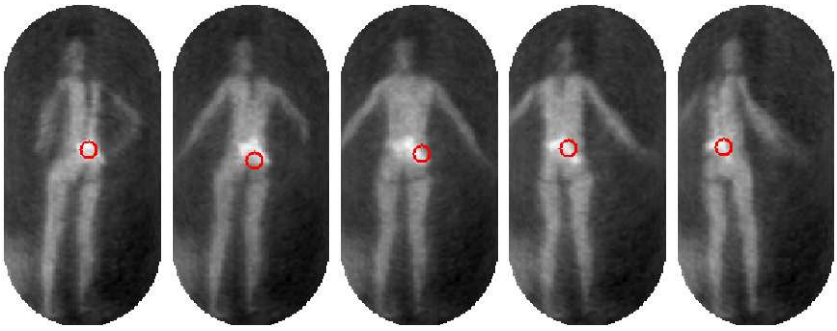

Fig. 7. Example of RPF tracking a target over the short period of time during which the target is both visible and well illuminated.

5.2 Computational Performance

For the results presented in this paper, the average processing time per frame was 0.83 seconds per frame on an Intel Pentium 4 (2.66GHz) running Fedora Core. This breaks down to 0.79 seconds for EM classification and 0.04 seconds for RPF filter tracking. It is likely that significant performance gains could be made with k-Means classification as proposed in Section 3.3. The complete system was written in C++ but no attempt has been made to optimise the computational performance of the code. Real-time operation of this system, with the MMW imager working at 12 fps, would be possible given suitable optimisation of the software and selection of an appropriate hardware platform.

6 Conclusions

We have presented a novel system for the automatic detection and tracking of metallic objects concealed under clothes using MMW sequences. The recent emergence of MMW video sensors makes our work very timely. To the best of our knowledge, no previous system combining MMW video imaging and advanced image processing techniques has been reported to date. Results have proven reliable on the current data test set. Further work that would

be required for this system to be operational would include a real-time implementation and extensive field-trials to determine accuracy, reliability and robustness in full scale operation. The requirement for real-time operation has been considered in this paper and work presented that would allow real-time operation given a suitable implementation on an appropriate hardware platform.

The long-term aim of this work is to develop a MMW system to enhance the public security screening process, such as airport passenger screening. This will involve work on detecting and characterising a wider range of threat materials, extending the range of tracking scenarios and incorporating human body models to improve tracking and provide 3-D visualisation preserving privacy.

Acknowledgements

We would like to thank Beatriz Grafulla González and QinetiQ for providing the data sequences and for technical discussions. This work is supported by EPSRC Research Grant GRS/68088 "ATRIUM" (Automatic Threat Recognition and Identification Using MMW).

References

Arulampalam, M. S., Maskell, S., Gordon, N., Clapp, T., February 2001. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. IEEE Transactions on Signal Processing 50 (2), 174–188.

Bar-Shalom, Y., Fortmann, T. E., 1988. Tracking and Data Association. Academic Press.

Bilmes, J., 1997. A gentle tutorial on the em algorithm and its application to parameter estimation for gaussian mixture and hidden markov models. Tech. rep., University of Berkeley.

Comaniciu, D., Meer, P., May 2002. Mean shift: A robust approach toward feature space analysis. IEEE Transaction on Pattern Analysis and Machine Intelligence 24 (5), 603–619.

Comaniciu, D., Ramesh, V., Meer, P., May 2003. Kernel-based object tracking. IEEE Transaction on Pattern Analysis and Machine Intelligence 25 (5), 564–577.

Coward, P., Appleby, R., August 2003. Development of an illumination chamber for indoor millimetre-wave imaging. In: Appleby, R., Wikner, D. A., Trebits, R., Kurtz, J. L. (Eds.), Passive Millimeter-wave Imaging Technology VI and Radar Sensor Technology VII. Vol. 5077 of Proceedings of SPIE. SPIE, pp. 54–61.

Grafulla-Gonzalez, B., Tomsin, M., Lebart, K., Harvey, A. R., June 2005. Modelling of millimetre-wave personnel scanners for automated detection. In: Proceedings of the International Symposium on Imaging for Crime Detection and Prevention,. IEE, London, pp. 9–13.

Haworth, C. D., González, B. G., Tomsin, M., Appleby, R., Coward, P., Harvey, A., Lebart, K., Petillot, Y., Trucco, E., 2004. Image analysis for object detection in millimetre-wave images. In: Appleby, R., Chamberlain, J. M., Krapels, K. A. (Eds.), Passive Millimetre-wave and Terahertz Imaging and Technology. Vol. 5619. SPIE, pp. 117–129.

Haworth, C. D., Saint-Pern, Y. D., Trucco, E., Petillot, Y. R., June 2005. Public security screening for metallic objects with millimetre-wave images. In: Proceedings of the International Symposium on Imaging for Crime Detection and Prevention,. IEE, London, pp. 1–4.

Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., Wu, A. Y., July 2002. An efficient k-means clustering algorithm: analysis and implementation. IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (7), 881–892.

Keller, P., McMakin, D. L., Sheen, D. M., McKinnon, A. D., Summet, J. W., February 2000. Privacy algorithm for cylindrical holographic weapons surveillance system. Aerospace and Electronic Systems Magazine 15 (2), 17–24.

Murphy, K. S. J., Appleby, R., Sinclair, G., McClumpha, A., Tatlock, K., Doney, R., Hutcheson, I., October 2002. Millimeter wave aviation security scanner. In: Proceedings of the 36th International Carnahan Conference on Security Technology. IEEE, London, UK, pp. 162–166.

Rabiner, L. R., February 1989. A tutorial on Hidden Markov Models and selected applications in speech recognition. Proceedings of the IEEE 77 (2), 257–285.

Sinclair, G. N., Anderton, R. N., Appleby, R., October 2001. Outdoor passive millimetre wave security screening. In: Proceedings of the 35th International Carnahan Conference on Security Technology. IEEE, London, UK, pp. 172–179.

Slamani, M.-A., Ferris Jr., D. D., 2001. Shape-descriptor-based detection of concealed weapons in millimeter-wave data. In: Casasent, D. P., Chao, T.-H. (Eds.), Optical Pattern Recognition XII. Vol. 4387 of Proceedings of SPIE. SPIE, pp. 176–185.

Slamani, M.-A., Vershney, P. K., Rao, R. M., Alford, M. G., Ferris, D., October 1999. Image processing tools for the enhancement of concealed weapon detection. In: Proceedings of the International Conference on Image Processing. Vol. 3. IEEE, Kobe, Japan, pp. 518–522.

Trucco, E., Plakas, K., 2005. Video tracking: a concise survey. IEEE Journal of Oceanic Engineering 30, (to appear).

Varshney, P. K., Chen, H.-M., Ramac, L. C., Uner, M., Ferris, D., Alford, M., October 1999. Registration and fusion of infrared and millimeter wave images for concealed weapon detection. In: Proceedings of the International

Conference on Image Processing. Vol. 3. IEEE, Kobe, Japan, pp. 532–536.
Xue, Z., Blum, R. S., July 2003. Concealed weapon detection using color image fusion. In: Proceedings of the 6th International Conference on Information Fusion. Vol. 1. IEEE, Queensland, Australia, pp. 622–627.