An Improvement of Subject Reacquisition by Reasoning and Revision

Jianbing Ma[1,2], Weiru Liu, Paul Miller, and Fabian Campbell-West

[1]School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast, Belfast BT7 1NN, UK
[2]School of Design, Engineering and Computing, Bournemouth University, BH12 5BB, UK
{w.liu,p.miller, f.h.campbell-west}@qub.ac.uk,
jma@bournemouth.ac.uk

Abstract. CCTV systems are broadly deployed in the present world. Despite this, the impact on anti-social and criminal behaviour has been minimal. Subject reacquisition is a fundamental task to ensure in-time reaction for intelligent surveillance. However, traditional reacquisition based on face recognition is not scalable, hence in this paper we use reasoning techniques to reduce the computational effort which deploys the time-of-flight information between interested zones such as airport security corridors. Also, to improve accuracy of reacquisition, we introduce the idea of revision as a method of post-processing. We demonstrate the significance and usefulness of our framework with an experiment which shows much less computational effort and better accuracy.

Keywords Subject Reacquisition, Time-of-Flight, CCTV Surveillance, Event Reasoning, Revision.

1 Introduction

During the last decade, there has been massive investment in Closed-Circuit TeleVision (CCTV) technology in the UK. Currently, there are approximately four million CCTV cameras operationally deployed. Despite this, the impact on anti-social and criminal behaviour has been minimal. Although most incidents, also called events, are captured on video, there is no response because very little of the data is actively analysed in real-time. Consequently, CCTV operates in a passive mode, simply collecting enormous volumes of video data. For this technology to be effective, CCTV has to become active by alerting security analysts in real-time so that they can stop or prevent the undesirable behaviour. Such a quantum leap in capability will greatly increase the likelihood of offenders being caught, a major factor in crime prevention.

To ensure in-time reaction for intelligent surveillance, one fundamental task to utilize CCTV videos is subject reacquisition[1] [1, 2, 28]. Subject reacquisition is the process of identifying a particular subject at a specific point in space and time given knowledge of a previous observation. There are many methods of performing reacquisition, but

[1] It is also called object reacquisition, here we use the term subject reacquisition since we focus on reacquiring people.

the general case is illustrated in Fig. 1. At each checkpoint (a camera/sensor) the face information is captured and stored in a central repository. A comparison of these features will reveal the most suitable match from a previous checkpoint for a subject at the current checkpoint.

In the context of this paper and the concept demonstrator it is based on, the sensors are standard CCTV cameras. A video analytic sub-system performs face detection and recognition. As the number of subjects in the system increases, the reacquisition problem becomes more difficult and mismatch becomes more likely. In addition, since comparisons are made between the live subject and all previously observed subjects the system is not scalable with large numbers of subjects. It is necessary to keep the potential number of comparisons at a manageable number, which is a primary goal of the event reasoning part.

To study subject reacquisition, in this paper, we consider a scenario which involves a simple secure corridor. At each end of the corridor is video camera. One agent is responsible for managing the events generated by these sensors. The aim of the system is to perform subject reacquisition. When presented with live data of subject X at the second sensor, the subject reacquisition problem is to identify which subject previously seen at the first sensor matches subject X (Same person always presents slightly different images in different sensors). In this context, subject reacquisition is also known as closed-set matching. This is one of the simplest scenarios that can benefit from artificial intelligence techniques. Also, note that this kind of secure corridor does exist in a set of airports, e.g., the Gatwick Airport[2].

In this scenario, two simplification assumptions are introduced and listed as follows:

- Only one subject shall be considered at one time(no occlusions etc.)
- The subject will be co-operative at key points (i.e. looking at the camera for a few frames)

The first assumption removes the need for tracking the subjects spatially and the second removes the risk that no face is present. These two assumptions are reasonable. Actually they are fully achieved in the secure corridor case in the Gatwick airport and other airports.

Currently, common approaches for subject reacquisition are simply by face recognition. Although face provides a rich source of information, its recognition comes with the price of uncertainty. A great deal of research effort has been applied to this area aiming for increasing robustness. In all non-trivial scenarios, there will be sources of errors in a pure video-analysis system and, if unchecked, these errors will propagate and cause further mistakes.

In this paper, an alternative approach is proposed where classic imperfect subject reacquisition (by face recognition algorithms) results are enhanced by using artificial intelligence approaches including event reasoning[3] and belief revision[4]. A real-world experiment is introduced and the baseline results are presented. More precisely, in this

[2] Currently they check passengers manually.

[3] As in [24, 18, 20, 25, 5, 16, 21–23, 7], event reasoning uses domain information to combine basic events to get high-level information.

[4] Belief revision depicts the process that an agent revises its belief state upon receiving new information [13, 19, 8, 9, 15, 12, 14, 11, 10, 17].

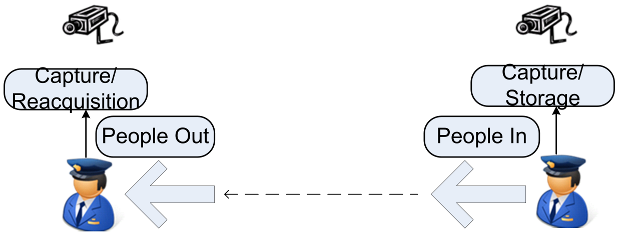

Fig. 1. Classic Subject Reacquisition

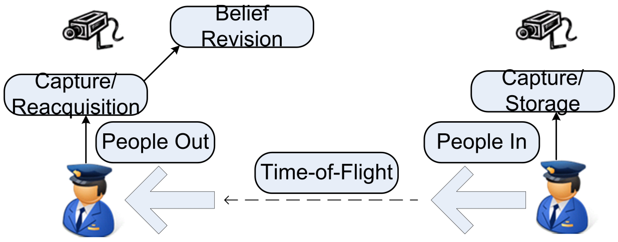

Fig. 2. Subject Reacquisition Enhanced by Reasoning and Revision

paper, event reasoning is used as pre-processing to reduce the size of comparison sets while belief revision, or the spirit of revision, is used as post-processing to enhance the recognition accuracy, as shown in Fig. 2. Here by comparison set we mean when a subject X is detected at sensor 2, the set of face models from sensor 1 that should be extracted to be compared with the face model of X.

The remainder of this paper is organized as follows. In Section 2, we briefly introduce our face recognition method. In Section 3, we state our event reasoning component using the time-of-flight information to reduce sizes of comparison sets. Section 4 describes the experimental scenario. In Section 5, we discuss how to use revision to improve accuracy of classification and reduce reacquisition failure. In Section 6, we conclude the paper.

2 Face Recognition

In this section, we only briefly introduce the method we used for face recognition since it is not the main focus of the paper.

The face recognition component is designed to be modular, so that different algorithms can be substituted for any component provided it has an appropriate interface.

With each new frame provided by the sensor it is necessary to search for faces. Face detection is a mature research area and there are many solutions available, among which we have tried two famous detectors: the Viola-Jones detector [27] and the Luxand detector [6]. The Viola-Jones detector is well-known, robust and good at finding faces in images in terms of precision and recall, but is not consistent in the bounding box returned of the detected face. So eventually we use the Luxand detector which provides more consistent results. Again, the Face Detection sub-system is modular, so any algorithm that can return a bounding box of a face in an image can be used.

We have also applied illumination compensation which is the process of adjusting the pixel intensities of an input image to match the illumination profile of a reference image. The algorithm used in the system is adapted from [4].

Faces of the same subject within a series of frames are registered to one face model and this model is stored in a buffer.

When a subject X is detected at sensor 2 reacquisition is performed by comparing every available face image to the set of known models (and corresponding subjects). And the subject in the set whose face model matches the most is chosen as the reacquired subject.

For each detected subject, a SubjectDetected event is created which contains four pieces of information needed. Let event i be denoted by the tuple

$$e^i = (o, d, t, M)$$

where o is the source identifier, d is the subject identifier, t is the timestamp and M is the face model. The source identifier is simply an integer that uniquely identifies the sensor. The subject identifier is also an integer, but is relative to the source. For example the subject with ID 1 at sensor 1 is not necessarily the same as the subject with ID 1 at sensor 2. The timestamp is the frame number, but it is assumed that this is synchronised across all sensors. The face model is a MACE filter produced from the face image. In implementation, M can be recorded by a pointer to the face model.

In addition, it is useful to discuss which methods can be used to integrate the reliabilites of face recognition generated by different frames. Now we use a simple maximum coverage method. That is, if a face model is considered the most plausible by face recognition in one frame, then its weight is added by one. After considering all frames, we normalize the weights of all possible face models. We are trying to use probabilistic and evidential [3, 26] methods to compare with the current one.

3 Reasoning by Time-of-Flight

When a subject is detected at sensor 2 the agent chooses face models from sensor 1 to match to the live data at sensor 2. Without limiting the comparisons the system is not scalable. In addition, the greater the size of comparison set the more likely it is a mismatch will be made. However, it is possible that by limiting the comparisons that the real subject is not included.

Our system uses domain knowledge, in the form of an average time-of-flight between the sensors, to determine appropriate models. For instance, a person cannot appear within the detection zone A, and then appear in zone B within 10 seconds. But if the time interval is changed to 10-20 seconds, then it could be possible; if the time-of-flight is 20-40 seconds, it is highly possible; if the time is greater than 40 seconds, it is less likely, etc.. Empirical probabilities for such information can be found by tests. However, by now, we will use simple settings that in some special range of time-of-flight, the certainty is 1 instead of probabilities.

Let E be the set of all events an agent has received so far. The most recent event is e^N (the N-th event). If $e^N.o = 2$, then the subject in this event needs to be reacquired. Let C be the comparison set of events that contain face models relative to e^N. The set C is composed according to the following rule:

$$C = \{e^i : e^i.o = 1, \tau_1 < e^N.t - e^i.t < \tau_2\}$$

where τ_1, τ_2 are thresholds on time-of-flight determined through domain knowledge indicating the lower and upper time limits for selecting subjects for comparison. Note that since e^N is the most recent event, $e^N.t \geq e^i.t$ is always true. This rule states that the time difference between the event of the current subject being detected at sensor 2 and the subject in event e^i being detected at sensor 1 is within the range (τ_1, τ_2). This rule gives event correlation between the SubjectDetected events at sensors 1 and 2. This is directly analogous to the data association problem in conventional tracking in which observations, i.e. SubjectDetected events, are associated with predictions, i.e. the time difference (τ_1, τ_2).

The average time-of-flight between sensors 1 and 2 is denoted by λ_F and is part of the domain knowledge of the application. This domain knowledge is used to calculate the time thresholds:

$$\tau_1 = \lambda_F - s\sqrt{\lambda_F},\ \tau_2 = \lambda_F + s\sqrt{\lambda_F}$$

where s in the two above equations is a scalar which determines how many standard deviations from the mean the threshold should be set.

4 Example Scenario

In the example scenario, 30 subjects pass between two checkpoints (sensors), as shown in Fig. 3. The time of arrival of each subject at a checkpoint is determined by a random

Fig. 3. All 30 subjects in the experiment

variable with a Poisson distribution. The Poisson distribution is used in queuing models for this purpose. It is characterized by a mean value, λ, which defines the probability distribution according to the standard Poisson distribution such that:

$$p(t; \lambda) = \frac{\lambda^t e^{-\lambda}}{t!}$$

where t is a non-negative integer.

Let the time of arrival of subject i at checkpoint 1 be denoted by t_i. Without loss of generality we set $t_1 = 1$. The difference in arrival time, $a_{i,j}$, between two consecutive subjects i and j can be modeled as a random variable with an Exponential distribution of mean denoted by λ_A. The arrival time of subject i at checkpoint 1 can be written as

$$t_i = t_{i-1} + a_{i-1,i}$$

for $i > 1$. For this scenario the values of $a_{i,j}$ were determined by a random number generator with $\lambda_A = 3$.

The time-of-flight between checkpoints 1 and 2 can also be modeled as a random variable with a Poisson distribution of mean λ_F. Let the time-of-flight for subject i be denoted by F_i. Then the time of arrival of subject i at checkpoint 2, denoted by u_i, can be written as

$$u_i = t_i + F_i.$$

For this scenario the values of F_i were determined by a random number generator with $\lambda_F = 11$.

The video data for the example scenario was created by manually editing existing video footage to meet the requirements of the timing model above.

The method for choosing the comparison set, C, is described in the previous section. Given the events produced for the example scenario, it is necessary to set $\lambda_F = 11$ and $s = 3.5$ to ensure the comparison set always includes the correct subject. At $s \leq 3$ the comparison set for some subject at sensor 2 will be empty.

For illustration simplicity, here we list the system behavior on nine subjects, which generates nine events at each sensor for a total of 18 input events. Reacquisition is performed for every input event at the second sensor for a total of nine output events.

Fig. 4[5] shows the comparison set for each subject. Fig. 4 provides the orders each

Fig. 4. The subject order at each sensor and the comparison sets for each subject

subject is present at each camera and the comparison set for each subject. To avoid

[5] For the sake of privacy, we have blurred the faces used.

confusion, the subjects are numbered according to their order at sensor 1. Between sensors 1 and 2 subjects 3 and 5 changed places. This means subject 5 moved quickly between the sensors and subject 3 moved slowly. Subjects 6 and 7 also changed places while moving between sensors.

Also from Fig. 4, we can see that our event reasoning method indeed reduces the size of comparison set a lot. Classically each subject at sensor 2 should compare to 9 nine candidates, but in our system, it only needs to compare to 5.3 candidates on average. That is, we have a 41% decrease in computational effort. Actually, in our 30 people experiment, the average size of comparison sets is 8, which saves 72% computational effort. Not surprisingly, if the scale of scenario becomes larger, it will save even more computational effort.

5 Revision

In this section, we discuss how to revise the face recognition results provided that there are obvious errors in the results. First, we describe the experimental result.

5.1 Comparison Reasoning

The subject reacquisition result based on the comparison set is shown in the following Fig. 5. In Fig. 5, we can see that in terms of classification results three out of nine

Fig. 5. Reacquisition result based on the comparison set

reacquisitions are incorrect, i.e., reacquisitions of subjects 1, 6, and 7, and one subject was not reacquired at all.

Classification accuracy is defined as:

$$\text{Accuracy} = \frac{\text{correct classifications}}{\text{total classifications}}.$$

The results above yield an accuracy of 67%.

Reacquisition failure can occur for the following reasons:

1. A subject is falsely detected and therefore the reacquired model is incorrect.
2. A subject is not detected when present at the second sensor, so reacquisition is not performed.
3. The subject is detected but wrongly reacquired. This can occur for the following reasons:
   (a) The comparison set C is empty.
   (b) The comparison set C does not contain the true subject.
   (c) The true subject is not chosen from the comparison set C.

Reasons 1, 2 and 3c on the list are failures of the video analytics, rather than the event reasoning part. Reasons 3a and 3b are failures of the reasoning.

Reacquisition failure is clearly shown in Fig. 5.1. Until a better description of sys-

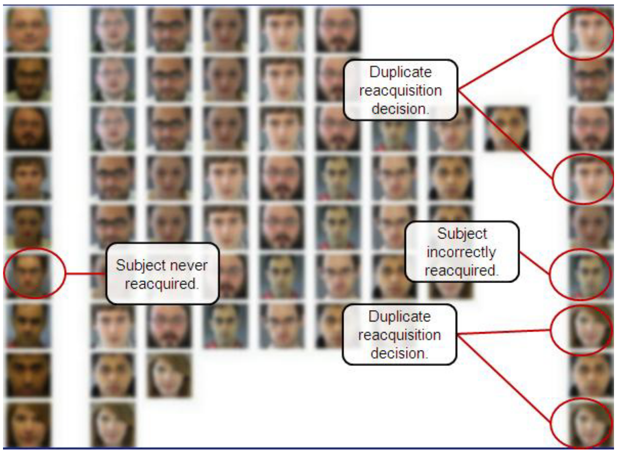

Fig. 6. Types of Errors in Reacquisition

tem failure is developed, the Reacquisition Failure measure will be calculated as

$$\text{Reacq.Failure} = \frac{\text{incorrect reacquisitions+missed reacquisitions}}{\text{possible correct reacquisitions}} = 4/9 = 44\%.$$

This is different to 1-Accuracy, as it considers models that are not reacquired at all. The failure measure is a worst-case view of the system, relative to the subjects rather than just the system outcome. It can be described as the proportion of 'problematic' subjects. Note that it is possible for the failure measure to be greater than 1, which reflects the imbalance between the number of ways the system can be correct versus the number of ways the system can fail.

Note that Subject 9 is incorrectly reacquired as the seventh subject at sensor 2. This raises the important issue of what to do once reacquisition has occurred. If the reacquisition is assumed correct, then subject 9 would be removed from subsequent comparisons and the final comparison set would be empty.

Note that if the system makes an incorrect classification that implies the system has, or will, either make a duplicate reacquisition decision or will switch subjects. A duplicate reacquisition implies that one of the subjects has not been reacquired. Subject switching is when two subjects are confused with each other.

In addition, in our 30 people experiment, the accuracy rate is 87% (26/30).

5.2 Revision of Reacquisition Results

When we have conflict reacquisition results, i.e., duplicate reacquisition, missing someone out, etc., we then use revision as a post-processing method to improve accuracy.

First, let us note that due to the closed world assumption, mistakes cannot happen in isolation, one mistake should infer another. That is, if we have missed out on someone, then there should be duplicate reacquisition, and vice versa.

The basic idea of revision in subject reacquisition can be described as follows: when we have conflict classification results, we can first determine the more reliable classification result (which is provided by the degree of certainties of face classification results), then we remove that candidate in the other conflicting comparison set. For instance, if both subjects X and Y are classified as person A while reliability (i.e., degree of certainty) shows the classification of X to A is more plausible, then we will remove candidate A in the comparison set of Y, and then we can choose the best match for Y in the new comparison set. This process is not strictly belief revision but deploys the idea of revision in artificial intelligence.

Here we should note that mistakes can cascade. That is, if in the revision process, we have wrongly removed a correct candidate from a comparison set, e.g., removing A from the comparison set of Y, and we choose B as the reacquisition result of Y, then it is possible that B has been reacquired by subject Z which makes another conflict between the reacquisitions of Y and Z and a further revision should be taken, and so on. In this sense, a wrong revision can destroy all correct reacquisition results.

To overcome this deficiency, we need to limit the amount of changes that can be made. Here we propose two kinds of revision:

- One step revision
- Limited revision based on threshold

By one step revision, we mean we do not proceed if a revision result induces a further conflict. For instance, suppose X and Y are classified as A, and Z is classified as B. Assume $r(X \to A) > r(Y \to A)$ (here $r(X \to A)$ is the degree of certainty for X classified as A), Y should be revised as a second choice, and suppose it is B which leads to further conflict with the reacquisition of Z. Now if $r(X \to A) < r(Z \to B)$, we just keep the original result, i.e., X, Y are classified as A and Z as B, else we keep the revised result, i.e., X is classified as A and Y, Z as B. And we finish here.

In Algorithm 1, notation C_Y means the comparison set of Y. By applying Algorithm 1, we have the following result showing in Fig. 7: From Fig. 7, we find that two incorrect reacquisitions have been corrected, which improves the Accuracy from 67% to 89%, a

Algorithm 1 One Step Revision

Require: All subjects, their comparison sets and their reacquired results.
Ensure: Revised reacquired results.

for each set S of conflict subjects **do**
    Revision = 1;
    A = classified result of each subject in S;
    $X = max_S\{r(X) : X \in S\}$;
    for each subject Y in $S \setminus \{X\}$ **do**
        $TempReacq(Y) = max_{C_Y}\{r(Y \to B) : B \in C_Y \setminus \{A\}\}$;
        B = TempReacq(Y);
        if exist Z, Reacq(Z)=B and $r(X \to A) < r(Z \to B)$ **then**
            Revision = 0;
            BREAK;
        end if
    end for
    if Revision = 1 **then**
        for each subject Y in $S \setminus \{X\}$ **do**
            $Reacq(Y) = max_{C_Y}\{r(Y \to B) : B \in C_Y \setminus \{A\}\}$;
        end for
    end if
end for
return Revised reacquired results.

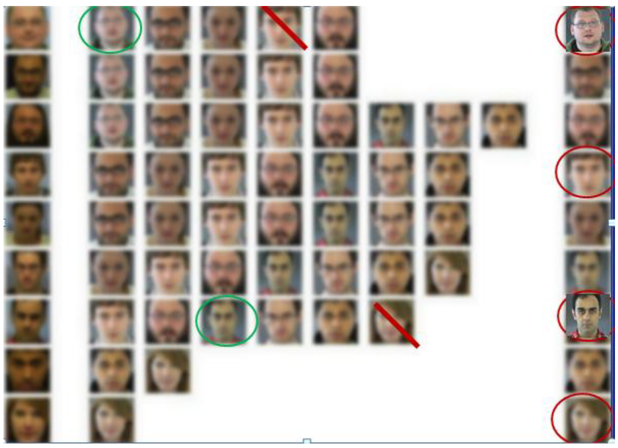

Fig. 7. Revised Reacquisition Result by One Step Revision

22% increase in Accuracy, and the reacquisition failure rate is reduced from 44% to 22%. Also, from Fig. 7, we know that revision cannot always achieve consistency.

In addition, for the 30 people experiment, the revised result achieves a remarkable 97% accuracy (29/30).

For limited revision, we should first introduce a threshold value t to indicate the difference between reliabilities (degrees of certainty) of classification results. As mentioned above, again we suppose X and Y are classified as A and Z classified as B.

If $r(X \to A) - r(Y \to A) >= t$, then Y should be revised to a second choice, else we do nothing. Now suppose Y is revised to B. Now

If $r(Y \to B) > r(Z \to B)$

if $r(Y \to B) - r(Z \to B) > t$, then we go on to revise the classification of Z, and so on.

else we do nothing.

If $r(Y \to B) <= r(Z \to B)$

if $r(Z \to B) - r(Y \to B) >= t$, then we go on to revise the classification of Y, and so on.

else we do nothing.

Each time reacquisition of some subject is changed, it may cause further revision. Since limited revision does not *rewind* as done in one step revision (e.g., the revision of Y from A to B can be rewound (Y changes back to A) by a further conflict reacquisition Z), the algorithm of limited revision looks very simple.

Algorithm 2 Limited Revision

Require: All subjects, their comparison sets and their reacquired results, a threshold value t, $0 < t < 1$.
Ensure: Revised reacquired results.
while exist a set S of conflict subjects **do**
  A = classified result of each subject in S;
  $X = max_S\{r(X) : X \in S\}$;
  for each subject Y in $S \setminus \{X\}$ **do**
    if $r(X \to A) - r(Y \to A) >= t$ **then**
      $Reacq(Y) = max_{C_Y}\{r(Y \to B) : B \in C_Y \setminus \{A\}\}$;
    end if
  end for
end while
return Revised reacquired results.

In this way, we can prove that revision (change of reacquisition result) can only happen for limited times, so we do not need to worry about the cascade mistake problem. We have the following result.

Proposition 1 *In Algorithm 2, revision for each subject can happen at most* $\lfloor \frac{1}{t} \rfloor$ *times. Here* $\lfloor \frac{1}{t} \rfloor$ *is an integer less than or equal to* $\frac{1}{t}$.

Proof of Proposition 1: We only need to note that each time the reacquisition for a subject is revised, its reliability to its reacquired result is reduced by at least t. So if it is revised by l times, then we have $lt <= 1$, which leads to $l \leq \lfloor \frac{1}{t} \rfloor$ (since l is an integer).

Of course, this $\lfloor \frac{1}{t} \rfloor$ is a theoretic limit. In practice, revision will happen much less. In our experiment, limited revision happens to get the same result as one-step revision, i.e., 97% (29/30). However, it could be expected that the results will be different when the experiment scale gets larger.

6 Conclusion

In this paper, we proposed a system for monitoring subjects passing between two sensors in a secure corridor. A hybrid real-synthetic scenario was created to model an authentic flow of subjects through the corridor. The sensors use video analytics to detect and learn face-based appearance models. Event reasoning using time-of-flight domain knowledge is proved helpful in reducing the comparison set and saves much computational effort. In addition, revision is deployed in this system to improve accuracy. Experimental results from the sequence show that revision does play an important role in increasing accuracy and decreasing reacquisition failure.

Applying artificial intelligence ideas to video surveillance is not a new idea, e.g., [24, 20], but it is seldom to see AI applied to subject reacquisition. Our paper hence can reminder researchers from the vision community and the AI community to be aware of the advantages of each area..

For future work, since the video analytics system is modular, there is potential for modifications to all aspects of the system. An evaluation of the different modules needs to be completed to ensure the best parameters are being used. Other video features, like clothing models and hair colour analysis can be incorporated.

For reasoning, currently generating the comparison is based on the time-of-flight of subjects through the corridor. The choice of the thresholds τ_1, τ_2 is most important. If the thresholds are too far apart the comparison set will include unnecessary samples. If the thresholds are too close the true subject may not be in the comparison set and the system will fail.

It may be more appropriate to use different values of s when calculating τ_1 and τ_2. A better scheme might be to start with strict bounds and relax them if the set C is empty. A problem with this, however, is that a non-empty set C that doesn't contain the true subject will definitely produce an incorrect reacquisition result. As a future work, we will investigate how to choose these two thresholds in general, and, if not well chosen, how does the system behave.

Detecting missed and/or duplicate reacquisitions is straightforward. Missed reacquisitions can be triggered by setting a maximum value on the time-of-flight between sensors. If a subject from sensor 1 has not been reacquired before that time then the agent must re-evaluate previous decisions with belief revision. Duplicate reacquisition occurs when two reacquisitions link to the same subject ID at the previous sensor.

Also, we are performing a much bigger scale of experiment to validate our system, and considering using other features (clothing color, Radio Frequency IDentification (RFID), etc.) to further improve accuracy.

References

1. C. Arth, C. Leistner, and H. Bischof. Object reacquisition and tracking in large-scale smart camera networks. In *Procs. of 1st ACM/IEEE Distributed Smart Cameras, ICDSC'07*, pages 156–163, 2007.
2. F. Campbell-West, H. Wang, and P. Miller. Where is it? object reacquisition in surveillance video. In *Procs. of Machine Vision and Image Processing (IMVIP'08)*, pages 182–187, 2008.
3. A. P. Dempster. A generalization of Bayesian inference. *J. Roy. Statist. Soc.*, 30, Series B:205–247, 1968.
4. X. Jiang, P. Fan, I. Ravyse, H. Sahli, J. Huang, R. Zhao, and Y. Zhang. Perception-based lighting adjustment of image sequences. In *Procs. of ACCV*, pages 118–129, 2009.
5. W. Liu, P. Miller, J. Ma, and W. Yan. Challenges of distributed intelligent surveillance system with heterogenous information. In *Procs. of QRASA*, pages 69–74, Pasadena, California, 2009.
6. Luxand.com. face recognition.
7. J. Ma. Qualitative approach to Bayesian networks with multiple causes. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 42(2):382–391, 2012.
8. J. Ma, S. Benferhat, and W. Liu. Revising partial pre-orders with partial pre-orders: A unit-based revision framework, 2012.
9. J. Ma, S. Benferhat, and W. Liu. Revision over partial pre-orders: A postulational study. In *Procs. of SUM*, pages 219–232, 2012.
10. J. Ma and W. Liu. A general model for epistemic state revision using plausibility measures. In *Procs. of ECAI*, pages 356–360, 2008.
11. J. Ma and W. Liu. Modeling belief change on epistemic states. In *Procs. of FLAIRS*, 2009.
12. J. Ma and W. Liu. A framework for managing uncertain inputs: An axiomization of rewarding. *Int. J. Approx. Reasoning*, 52(7):917–934, 2011.
13. J. Ma, W. Liu, and S. Benferhat. A belief revision framework for revising epistemic states with partial epistemic states. In *Procs. of AAAI*, pages 333–338, 2010.
14. J. Ma, W. Liu, D. Dubois, and H. Prade. Revision rules in the theory of evidence. In *Procs. of ICTAI*, pages 295–302, 2010.
15. J. Ma, W. Liu, D. Dubois, and H. Prade. Bridging Jeffrey's rule, AGM revision and Dempster conditioning in the theory of evidence. *International Journal on Artificial Intelligence Tools*, 20(4):691–720, 2011.
16. J. Ma, W. Liu, and A. Hunter. Inducing probability distributions from knowledge bases with (in)dependence relations. In *Procs. of AAAI*, pages 339–344, 2010.
17. J. Ma, W. Liu, and A. Hunter. Modeling and reasoning with qualitative comparative clinical knowledge. *Int. J. Intell. Syst.*, 26(1):25–46, 2011.
18. J. Ma, W. Liu, and P. Miller. Event modelling and reasoning with uncertain information for distributed sensor networks. In *Procs. of SUM*, pages 236–249. Springer, 2010.
19. J. Ma, W. Liu, and P. Miller. Belief change with noisy sensing in the situation calculus. In *Procs. of UAI*, 2011.
20. J. Ma, W. Liu, and P. Miller. Handling sequential observations in intelligent surveillance. In *Proceedings of SUM*, pages 547–560, 2011.
21. J. Ma, W. Liu, and P. Miller. A characteristic function approach to inconsistency measures for knowledge bases. In *Procs. of SUM*, pages 473–485, 2012.
22. J. Ma, W. Liu, and P. Miller. Evidential fusion for gender profiling. In *Procs. of SUM*, pages 514–524, 2012.
23. J. Ma, W. Liu, and P. Miller. An evidential improvement for gender profiling. In *Procs. of Belief Functions*, pages 29–36, 2012.

24. J. Ma, W. Liu, P. Miller, and W. Yan. Event composition with imperfect information for bus surveillance. In *Procs. of AVSS*, pages 382–387. IEEE Press, 2009.
25. P. Miller, W. Liu, F. Fowler, H. Zhou, J. Shen, J. Ma, J. Zhang, W. Yan, K. McLaughlin, and S. Sezer. Intelligent sensor information system for public transport: To safely go .... In *Procs. of AVSS*, 2010.
26. G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
27. P. Viola and M.J. Jones. Rapid object detection using a boosted cascade of simple features. In *Procs. of IEEE CVPR*, pages 511–518, 2001.
28. M. R. Walter, Y. Friedman, M. Antone, and S. Teller. Appearance-based object reacquisition for mobile manipulation. In *Procs. of IEEE Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1–8, 2010.