Integrated roadmap for the rapid finding and tracking of people at large airports

Henri Bouma,* **Jeroen van Rest, Kim van Buul-Besseling, Jacomien de Jong, Anton Havekes**

TNO, Oude Waalsdorperweg 63, 2597 AK The Hague, The Netherlands

Received: February 4, 2014; Revised: April 20, 2015; Published online: 19 November 2015

Abstract

This paper presents an integrated roadmap for the rapid finding and tracking of people at large airports. Business continuity at an airport is threatened by incidents such as left luggage and breaches of secure areas. These activities require rapid response because they are high frequency, high risk and high cost. For these reasons, live tracking and forensic searches of people are important tasks, for example, finding the owner of a left piece of luggage or tracking a trespasser. Finding and tracking people using surveillance cameras without technical aids is time consuming. The roadmap for the rapid finding and tracking of people at airports is based on discussions with end users, an assessment of the state-of-the-art and an integral assessment of work processes, human-machine interfaces, computer vision and the information and communications infrastructure. According to the roadmap, a major step forward can be achieved by integrating advancements in four areas: (i) increasing camera coverage; (ii) implementing computer vision for automated recognition and tracking; (iii) using contextual user interfaces; and (iv) adapting training and work processes to the tasks of finding and tracking people.

Keywords: Airport; Surveillance; People; Forensic Searches; Tracking; CCTVs; Roadmap

1. Introduction

Large airports are complex critical infrastructures that involve high stakes for the national economy and the safety of people working, visiting and transiting the airport. Airports formulate strategic business goals in terms of the number of passengers, airport reputation, security and personnel costs, etc. These business goals can be ultimately translated to the need to control and mitigate the risk factors that negatively impact airport business continuity and reputation.

The security of an airport is constantly affected by incidents that could potentially grow to be a threat to business continuity (e.g., a large-scale evacuation). Examples of frequent incidents are left luggage, trespassing in secure areas, pickpocketing, shoplifting, swindling, drug dealing and harassment of airport personnel or members of the public. Less frequent, but more significant, incidents include hijackings, bombings, shootings, robberies and terrorist acts. When an incident causes an evacuation or a perimeter to be established that includes airline gates, the continuity of air traffic is affected and the chain of events potentially impacts business continuity at other airports. It may take days, sometimes even weeks, before all the delayed passengers and their luggage are handled. The actual costs depend on factors such as flight frequency, airlines, presence of alternatives and percentage of affected flights. A rough estimate is that a single incident may cost millions of euros distributed over multiple companies and it may damage the reputations of the airport, carriers and air travel, in general. Apart from the less frequent large-scale incidents, frequently-occurring small incidents impose significant structural pressure on security departments at airports.

* E-mail: henri.bouma@tno.nl; phone: +31 888 66 4054

The final publication is available at Elsevier's ScienceDirect: http://dx.doi.org/10.1016/j.ijcip.2015.11.002

Airport security personnel dedicate a substantial portion of their time to searching and tracking people in order to prevent and respond to incidents. The more time it takes to find a person of interest after an incident is detected, the more complex the search becomes because the area to be searched increases with time. Many security organizations use closed-circuit TV (CCTV) systems to help assess, prevent, mitigate and respond to incidents. However, addressing incidents in an effective manner requires more than simply purchasing expensive equipment [16]. Policy makers, airport management and suppliers need an understanding of the investments required to accomplish the joint strategic goals. Therefore, it is important to understand the maturity of the technologies and if they are ready for operational use. Acquiring this understanding before the competition can give an airport significant competitive advantages.

This paper attempts to answer three questions. How can airport security at a large civilian airport find and track people faster? How can this contribute to strategic business goals? How mature are the available technologies and what technologies are feasible in the short and long terms? Researchers have developed automated techniques for tracking [4], recognition [3, 18] and other aspects of identifying persons of interest in CCTV footage. However, they do not provide an integrated overview of the various solutions for finding and tracking people at an airport and a plan that matches the maturity of the solutions with short-term and long-term goals. Researchers have also created roadmaps for generic airport security [14]. However, these roadmaps are typically at a high level of abstraction. The roadmaps may cover a wide range of topics, but the level of abstraction forces them to omit many aspects that are relevant to finding and tracking individuals.

The novel contribution of this paper is an integrated roadmap geared towards the rapid finding and tracking of people at airports while maintaining the same or higher levels of security. The roadmap focuses on work processes, human-machine interaction, computer vision technologies and the information and communications technology (ICT) infrastructure. The roadmap distinguishes between technologies that can be used within one year, those that can be expected in the near future (three years) and those that are anticipated in the long term (beyond five years).

2. Roadmap creation

A technology roadmap is a plan that matches short-term and long-term goals with specific technological solutions. The content of a roadmap is determined by four aspects: (i) starting point; (ii) desired end condition; (iii) external factors; and (iv) scope. Only after all these aspects are determined can the internal structure and content of the roadmap be determined. This section describes the methodology used to obtain the roadmap. Next, it describes the four aspects that determine the content of the roadmap and elaborates on the use cases. Finally, it presents the main components of the roadmap.

2.1. Methodology

Three steps were used to obtain the roadmap described in this paper. First, the frequency and impact of incidents were determined in collaboration with the Dutch Military Police and Schiphol Airport Security. This study resulted in the selection of the use cases described later in this section. Second, information was gathered about the organizational processes at Schiphol Airport by interviewing experienced security and airport services personnel about the most relevant use cases, the current working process, interfaces and technologies. Third, technology-related information was collected by consulting experts in the fields of information and communications technologies, computer vision technologies, human factors, and behavioral and social sciences. This enabled the creation of a survey of relevant research and state-of-the-art commercial products.

The usability of a technology depends on the performance of the technology and the circumstances in which it is used. Existing technology can be tested in realistic environments by users. However, the prediction of future technological developments requires input from specialists. A questionnaire was formulated to obtain assessments from computer vision experts to underpin the time estimates in the roadmap. This way, deep insights into the obvious and uncertain aspects of technological developments were obtained. Although the responses to the questionnaire may be biased (possibly, too optimistic or too pessimistic), it is clear that the relative maturity of technologies could be assessed reliably by the participating experts.

2.2. Aspects that determine roadmap content

The content of a roadmap is determined by four aspects: (i) starting point; (ii) desired end condition; (iii) external factors that may influence the roadmap; and (iv) scope of functionality that it should cover at certain points in time. Only after all these aspects are determined can the internal structure and actual content of the roadmap be determined.

The starting point in the roadmap is a fictional large airport that is similar to many existing airports. At the airport, several organizations and teams are involved in finding and tracking people, for example, border/military police, customs and airport personnel [20]. Multiple surveillance teams operate among the public without access to CCTV footage. Each organization has its own camera surveillance room with camera operators. When personnel in the field need to locate an individual, a call is made to the appropriate camera surveillance room with a search query comprising the person's description, time, place and/or behavior. Control and coordination for finding and tracking people are centralized. The user interface of the surveillance system is a video wall with multiple feeds. Computer vision technology that analyzes video content is not used. A secure fixed IP network is in place for transporting digital video data, but a secure mobile network is not available. About 80% of the airport floor is covered by CCTV systems – a large portion is covered by pan-tilt-zoom cameras with approximately 10% continuous coverage. Other organizations in other rooms also use the pan-tilt-zoom cameras so these cameras cannot always be used for finding and tracking purposes. All the entrances are covered by cameras, but frontal images of people as they enter and leave the airport are not taken and the resolution of facial images is limited. The cameras have some on-board processing capabilities, but these are not used for video content analysis.

The second aspect, the desired end condition, is that a person should be found within minutes or even seconds based on an image or description. People can be tracked with minimal human assistance. Surveillance and intervention teams work autonomously, even when they are mobile. Cooperation between different teams and organizations is not hindered by technology. Privacy and data protection risks are mitigated adequately. Key performance indicators are the time needed to find or track a person and the number of operators needed to perform the task.

The third aspect that determines the content of a roadmap is external factors. The roadmap is constructed under the assumption that, in the future, the terrorist threat level may be elevated but, in general, the threat level is low. Thus, the required resources should be low. Security awareness amongst the public is generally low, but it may be temporarily enhanced by incidents. This means that similar behavior by members of the public is anticipated as in the past (e.g., amount of left luggage and number of trespassers). Since the service level required to remain competitive with other modes of transportation will increase, use cases that are not security related (e.g., locating and tracking people who need assistance) must be included. Privacy awareness is expected to increase; this is influenced by stricter legislation and increased media attention on privacy incidents in general. The roadmap must, therefore, explicitly address ethical and privacy issues.

The fourth aspect is the scope of functionality. The scope is expressed in terms of the use cases it facilitates. In this work, the selection of use cases for airport surveillance is based on the frequency, potential impact, relevance to airports and advantages arising from identification and tracking technologies. Detailed data about the numbers and frequencies of incidents are usually confidential for security and business reasons and they may be inaccurate due to the use of different counting methods. Based on discussions with several airports, four use cases were selected due to their high frequency: (i) left luggage; (ii) breaches of secure areas; (iii) harassing persons; and (iv) slow passengers. The slow passengers use case is not a security threat, but an example of where providing improved service can help maintain a competitive position.

2.3. Use cases

The first use case involves left luggage, which is luggage that does not have an obvious owner in the vicinity. This occurs many times a day at large airports. Small bags or jackets may be inspected to identify their owners or to confirm that they do not contain explosives. However, most left luggage is not opened immediately because safe and easy visual inspection is not possible. In order to determine if an item of luggage contains explosives, a perimeter is established and a bomb sniffing dog or explosives detector is brought to the scene. In rare cases (still, multiple times per year), traces of explosives are detected. A larger perimeter is then established based on the size of the luggage and the bomb squad is summoned. In such a situation, business continuity is seriously affected – delayed flights, delayed passengers and blocked ground transportation.

In the vast majority of cases involving left luggage, there is no hostile intent. Locating the owner of the luggage quickly saves a considerable amount of time and effort. An innocent owner generally attempts to recover the lost item as soon as its loss is noticed. This may delay a particular flight, but it also provides a service opportunity for the airport to help a distressed client. In order to locate the owner, it is common to first examine the recorded video to obtain a description of the person who left the luggage and then track the person forward in the recorded video streams until the current location of the person is found. The behavior and social interactions of the person with other persons (before and after leaving the luggage) are analyzed to assess whether the person may have had hostile intent. This use case requires the rapid detection of left luggage, inspection of the person who left the luggage, finding the current location of the person and tracking the person until security personnel approach the person.

The second use case involves a breach of a secure area. A breach involves one or more persons who enter a secure area in an unauthorized manner such as using an emergency exit to go from landside to airside, or vice versa for illegal immigrants. When a breach occurs, security personnel must find the person quickly because, in the worst-case scenario, the intruder may be carrying prohibited material (e.g., explosives or weapons). The effect of a breach is that the entire area is no longer secure. If this unsecure area is airside, then the terminal has to be cleared, flights are cancelled or delayed, arriving passengers may not be allowed to disembark, transit passengers miss flights, airline crews have to be replaced and hotel accommodation has to be arranged for stranded passengers and crews. If the moment of the security breach is caught on video and seen directly at a surveillance center, then the trespasser can be found and apprehended in a short period of time. If sufficient cameras exist, contacts that the trespasser had with other people would also have been recorded. This may eliminate the need to evacuate the complete airside. This use case requires rapid detection of the breach, obtaining a description of the trespasser, finding the current location of the trespasser, tracking the movements of the trespasser and the people he contacted until security personnel can apprehend all the individuals.

The third use case involves a person who for some – typically, mental, emotional or medical – reason disturbs others. The use case begins when a complaint is received by public services or airport staff. By the time the complaint reaches security personnel, the person may have left the scene. If the problems involving the person are not addressed promptly, there is a good chance that he or she will harass others as well. The main issues in this situation are to understand the behavior of the person, identify the current location of the person and prevent the next incident involving the person. Such incidents cause considerable irritation and decrease public perceptions of safety at the airport. If the person is apprehended quickly, the source of negative effects is eliminated. Depending on the reason for the misbehavior, the person may have to be helped as well. This use case requires finding the current location of the person and tracking the person until security personnel arrive at the scene.

The fourth use case involves a passenger who causes a delay due to age (young or old), a handicap or poor understanding of airport signage. Airport shops are a large source of revenue, but they can slow passengers and cause delays. This is an opportunity to improve airport operations and the perceived service level, but only if early intervention is possible, which requires the localization of the passenger who is delaying a flight. Audio messages are often used to remind tardy passengers to get to their gates. The actual linking of a passenger to a departing flight requires identification, but after this is done, only regular recognition and tracking is necessary.

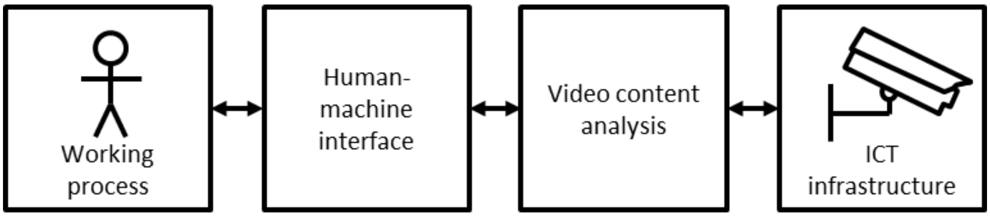

Figure 1: The four main elements of the roadmap for rapidly finding and tracking people at airports.

2.4. Roadmap overview

The roadmap has four main elements: (i) work processes; (ii) human-machine interfaces; (iii) video content analysis; and (iv) information and communications technology infrastructure. These elements were selected to cover the entire process from the camera to the human operator. Cameras generate images that must be transferred using the information and communications infrastructure. Video content analysis extracts high-level information from image pixels. Human-machine interfaces facilitate the connections between operators and the systems. The work processes address the various preconditions to establish and ensure that the complete chain is effective and accomplishes the business goals (Figure 1).

3. Work processes

International airports have strategic goals such as increasing business continuity, reducing personnel costs and increasing the service level. The work processes must be optimized to achieve these goals. This section focuses on three elements: (i) the environment around an employee (i.e., coordination and collaboration); (ii) the employee (i.e., training and performance evaluation); and (iii) relevant preconditions to perform the tasks (i.e., ethical and legal aspects).

3.1. Coordination and collaboration

At international airports, several interdisciplinary security teams and organizations are operational at the same time in the same area [13, 15]. As far as finding and tracking people at airports is concerned, many possibilities exist for improving coordination and collaboration to avoid cross-impacts. Finding and tracking people is a joint task performed by operators in surveillance rooms and mobile personnel in the field. Both groups of people need up-to-date operational pictures to execute their tasks, for which communications are essential. Communications are often in place between the organizations, but they can seek mutual support to improve their work processes (e.g., joint surveillance teams, lining-up work processes such as camera usage and centralized surveillance) [1]. In about three years, it should be possible to de-conflict potential inter-organizational overlaps by mutually providing more real-time insights into the consequences of planned actions by other organizations.

Some critical success factors can be identified to improve coordination and collaboration [5]. These include clear role definition and professional attitudes, open and direct communications, trust and a common culture. The success factors can be improved by collaboration awareness and shared training programs. Knowing and understanding the needs, goals, expectations, culture, capabilities and procedures of partners makes collaboration more effective [39].

3.2. Training and performance evaluation

Obviously, personnel must be well trained and educated to perform their tasks of finding and tracking people. They must learn the work processes and use novel technologies to improve their performance. Changes in the organizational structure have consequences, but they also offer opportunities for assessment, training and education. For example, new roles, tasks and responsibilities for personnel in an organization can be examined in the short term. This includes an analysis of the required competencies, which results in new training regimens, evaluations of performance and, possibly, the need to hire new personnel. Collaboration may include joint technical sessions for shared information and communications infrastructures. In joint training sessions, experiences, good practices and incident scenarios are shared, which contributes to trust and a common culture.

In the long term, inter-organizational training can be a force multiplier. Flexible capacity management can be enhanced by common capabilities and career paths. This can lead to more flexibility in exchanging personnel during periods of overload.

Finally, training and education are necessary to use finding and tracking techniques. Situational awareness is crucial to operators performing their tasks well. Endsley et al. [17] differentiate between three levels of situational awareness: (i) perception of data; (ii) comprehension of the significance of the situation; and (iii) projection of future states. Advanced technical support can improve awareness, for example, by providing temporal information that can support higher levels of awareness. Using surveillance techniques and reacting correctly require education, practice and experience. Training is necessary to ensure that information is perceived and projected as intended. Personnel who use technology indirectly also need to be informed about the power and risks of the surveillance techniques. This can be accomplished, for instance, by information sessions or real-world training.

3.3. Ethical and legal aspects

Society, businesses, organizations, technologies and threats change. Because of these changes, it is imperative to consider the ethical and legal aspects of implementing or changing work processes and technologies from the very beginning. The risks of not doing this properly are reduced effectiveness and potential loss of investment and societal support. The solutions described in this paper go beyond traditional surveillance and raise additional legal and ethical issues that cannot be covered by existing privacy measures.

For example, sending images to improve shared situational awareness carries security and data protection risks. Several proposed solutions (e.g., more cameras) result in an increase in the amount of stored personal data; this leads to a data protection issue that must be addressed properly. In addition, the effects of a data breach can increase when a finding and tracking system is in place. If the implementation stores tracks or (cut-out) images of people, then it is much easier to do harm with the data. The concrete impact of these issues cannot be reliably predicted without detailed designs.

Systems for finding and tracking people engender biases. A tracking system may encourage an operator to keep looking at a person who has performed a minor deviant act (case building and confirmation bias [32]). Also, a system that facilitates searching and filtering based on certain characteristics can lead to surveillance bias – the predisposition to focus on certain people more often than others regardless of their actual behavior (e.g., because of their ethnicity, skin color, sex or age). Moreover, a tracking system may be influenced by positive feedback, which is an effect that reinforces itself. A system that can track people with locally-unique characteristics (e.g., bright colored clothing) better than other people, tends to collect more images of locally-unique people. An increased number of images taken of a person can further amplify the effect and improve tracking because the person is imaged from more angles and under a greater variety of conditions. People who are locally unique due to their appearance may be easier to find than people who blend in; this may also lead to bias.

Privacy by design is a framework for mitigating privacy and data protection risks [40]. The framework incorporates methodological, organizational and technical measures that can be made concrete for specific application domains. Organizational measures may include legal agreements on exchanging information. Technical measures may involve obscuring faces [33] or using encrypted metadata databases [28]. An important component of the technical framework is a privacy impact assessment of surveillance [42]. It is recommended to perform such an assessment during every major design effort and to use the results in the next design iteration.

When the first new functionality becomes available, the concrete impact should be assessed. This will enable designers to make adaptations and will help manage surveillance processes in an appropriate manner.

4. Human-machine interfaces

A human-machine interface (HMI) is an interface through which an operator assesses the situation and controls the surveillance system (observe and act in the observe-orient-decide-act (OODA) loop [34]).

4.1. Situational awareness

Optimal situational awareness is crucial for a camera operator, especially because his/her tasks are invariably performed under time pressure. Following a person throughout an airport on camera requires knowledge and interpretation of the camera view in relation to the context and the environment (e.g., physical layout of the airport, knowing which camera to select and what action to take).

At most large airports, a person of interest is tracked over multiple cameras by an operator who determines the next camera where the person is expected to appear. The operator has to manually aim this pan-tilt-zoom camera. Also, when switching from one camera to another, the viewing direction is often not aimed at the preferred direction because of the pre-set viewing direction. This requires the operator to have a mental map of what is seen and how to control the camera to quickly regain sight of the person being followed. When a suspect moves rapidly, it is common to have a second operator help set up the next camera because of the risk of losing sight of the person.

The situational awareness of the camera operator can be increased by a geographical representation of the location, including the camera positions. Often, a static map with camera positions is already available. A first step is to use an interactive map (geographical information system) with camera locations, additional real-time viewing directions, zoom level indications, viewing height and field-of-view obstructions (available from S&T Corporation). Figure 2 shows camera views on an interactive map, which provides an intuitive representation of the relations between camera positions and field-of-view indications. Providing an operator with an intuitive representation of relations between camera positions reduces operator training time and speeds up the task of switching cameras to track a person. Mapping camera information to real-word coordinates is essential to creating an interactive map. This technology is already available in several products (e.g., from Genetec, Axxon and S&T Corporation [25]), but it is not commonly used at airports.

Additional context information in selectable layers can provide even more situational awareness (e.g., Google Earth, street view, walking routes and store information). Also, this presents dynamic information about hotspots, events and crowds.

Knowledge of the current positions of personnel in the field helps enhance operator situational awareness. Maintaining communications with field personnel and having “eyes on the floor” are also important for operators. Tracking field personnel and presenting their positions in an interactive map enables an operator to contact the right person during an incident, which results in a quicker response. Situational awareness can also be enhanced using a 3D representation (instead of a flat 2D map) upon which camera images are patched. However, in a complex environment like an airport with many camera perspectives, floors, staircases, corners and obstructions, creating an intuitive representation can be very challenging. Finding people is an important task for operators in a control room and for mobile personnel in the field. Search interfaces on mobile clients that respond to textual descriptions (e.g., from Genetec) may be used to enhance shared situational awareness and communications.

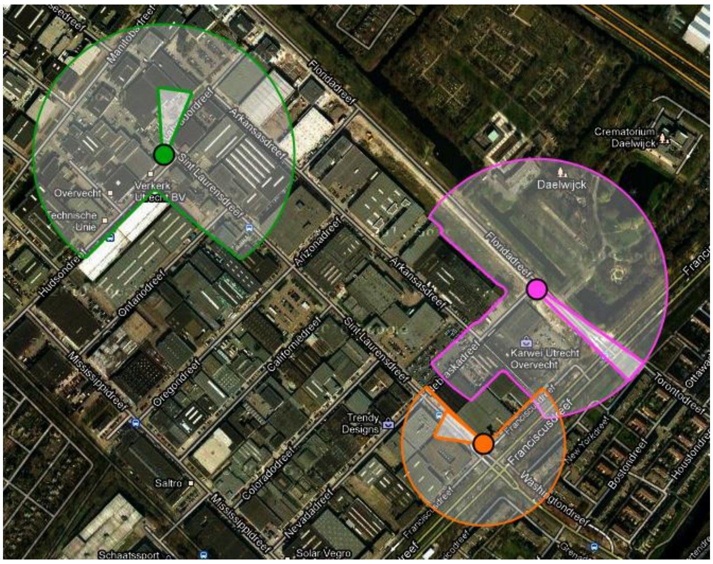

Figure 2: Camera views shown on an interactive map.

4.2. Surveillance system control and communications

Several control tasks can be distinguished: selecting and orienting a camera, entering a search query, sifting through search results and setting the focus of attention (selecting a person to be tracked). These control tasks can be performed efficiently and effectively via intuitive interactions between the operator and system. A first and easy step is to provide indicators or buttons inside a camera view that directly switch the view to the next camera. This reduces the cognitive load on the operator because he does not have to remember which camera has to be addressed next. These features are already available in commercially-available software. The next step is to utilize the interactive map (Figure 2) not only to enhance situational awareness, but also to select and control the cameras directly using the map. When an operator receives a notification that a person was seen at a specific location, clicking on the corresponding position on the map will automatically aim the surrounding cameras at the spot. It is even possible to take field-of view blocking into account (based on a height map) to ensure that a camera focuses on a viewing area instead of a blank wall. This way, the operator does not have to select, manipulate and focus the surrounding cameras sequentially, but is just one click away from directly searching multiple camera views for the person of interest.

Interaction with field personnel is a key task for an operator. Knowledge of the position and status of field personnel (e.g., busy, available or patrolling) is important, along with the ability to directly contact them by selecting them on the map; this makes the interactive map the central and intuitive human-machine interface for the camera operator and enables quick, direct and effective communications and task allocation. Based on knowledge of the occupations, task loads and positions of operators and personnel in the field, tasks can be adaptively allocated by the system or a supervisor. For example, in the case of a security breach, the nearest available field personnel could be automatically directed to the precise location by the system.

When performing the tracing and tracking tasks (e.g., for the left luggage use case), the operator searches for a person sequentially in historical camera images based on the known time, description and location. This can be very time consuming. However, the search space can be reduced by taking spatial and temporal constraints into account. Cooperation between the human and the system helps reduce false positive and false negative errors. Assistive technology can also be effective when a secure zone is breached. After a breach of a secure area, an alert is generated, cameras are automatically directed towards the exits and the related video streams are presented to camera operators and sent to field personnel. The benefits of this approach are direct identification and significant time savings.

New opportunities and challenges with regard to operator interaction are brought on by the use of (semi-) automated search and tracking systems. These systems significantly enhance operator efficiency and effectiveness, but challenges arise in formulating and activating searches, responding to false positives and interacting with the systems. While fully-automatic multi-camera tracking is not yet feasible, cooperation and/or smart dialog between the operator and the system are feasible. For example, an operator could select a location and the description of a person (e.g., wearing a red shirt) and the system would proceed to search rapidly through the saved images for about 20 candidates near the location who meet the description; the operator could then quickly identify the right person and notify the system. When tracking a person through an airport, the operator could mark the person in a camera view and let the system follow the person – as long as the person remains visible in the view – and even let the system suggest alternative cameras based on a predicted walking pattern.

5. Video content analysis

Several surveillance tasks can be (partly) automated, including detecting left luggage, detecting suspicious behavior and events, and searching for and tracking people. The operators are supported by computer-vision techniques that assist them in analyzing potentially large numbers of video images.

5.1. Left luggage detection

Several CCTV products are available for detecting left luggage [19, 26, 38], including systems from IPS-analytics, Aralia, AgentVI, Axxon, Bosch, Ipsotek, ioImage and ObjectVideo. Although CCTV technology is quite mature, it is generally not effective for detecting left luggage in large crowded environments. However, when an abandoned piece of baggage is discovered by other means, the technology is easily used to identify the time of placement with minimal human interaction. The principal advantage is that the time taken to identify the person who left the item of luggage is significantly reduced. Indeed, airports could immediately start pilot implementations involving human-machine collaboration for left luggage detection.

5.2. Behavior and event recognition

Behavior and action recognition is currently an active research topic [10, 11, 12] and some tools have demonstrated high accuracy in controlled environments [11]. However, the reliable detection of complex suspicious behavior in a realistic environment does not yet reach a performance level that is adequate for airport security. Nevertheless, simple events and actions can be detected quite reliably at selected locations for dedicated applications such as intrusion detection in a secure zone, waiting time in a queue, people passing a virtual tripwire, people entering and leaving a region, and loitering detection. Detecting a person who is approaching a closed perimeter could be valuable in the use case of a breach of a secure area.

5.3. Tracking and searching

Automatic single-camera systems for tracking people in sparsely-populated areas are readily available. Fully-automatic, multi-person tracking using multiple cameras with non-overlapping (disjoint) views in a crowded environment is currently not available as a reliable product for end users [23]. However, it has been shown that an interactive system can improve the performance of a human operator in such cases [8]. Another approach for improving tracking is to combine views from multiple angles [24]. This requires multiple CCTV cameras to cover a tracking area.

Several systems that claim to support "forensic searches" of video recordings could be used for surveillance purposes; these systems are manufactured by Aralia, AgentVI, 3VR, Nice, Ipsotek and Bosch. The products are still quite novel and are not used at many airports; moreover, some products are optimized for locating cars and other objects, not people. It is most common to attempt to recognize people using a query image [22, 36]. However, it is also becoming feasible to use a semantic description of a person's clothing [7, 37] (a feature provided by a product from Aralia Systems). A semantic description is especially useful when no example images are available to perform a search.

Three options are available to improve the finding and tracking of people starting with a current installed base: (i) execute one of the aforementioned software systems on a central cluster of computers (i.e., core); (ii) use basic intelligent video analytics functionality that is integrated in cameras (i.e., distributed and at the edge); and (iii) use a better recognition methodology in combination with intelligent video analytics integrated in cameras for higher quality computations at the cameras.

The first solution is to compute everything at the core. Several systems are available that claim to support forensic searches, including products from Aralia, AgentVI, 3VR and Ipsotek. The disadvantage of using software that is not integrated in a camera is that the computations must be done at the core. When limited processing capacity is available at the core, only a restricted small-scale search of video from a limited set of cameras is supported; fast interaction with many cameras is not feasible. In some cases, this may be the optimal solution, for example, finding one person in a limited area where a delay due to pre-processing is acceptable. However, the computational power must shift from one region to another as the query propagates through an airport to find and track the person of interest.

The second solution is to use existing cameras for detection and to compute metadata (e.g., histograms or more advanced descriptors [2, 6]). Most cameras at large airports support intelligent video analysis functionality. This allows the segmentation of moving objects, tracking using a single camera and searches. The advantage of this approach is that the computations are performed by the camera and that a huge amount of video data is transformed to metadata, which supports efficient searches and retrieval of information. This enables a large-scale search involving many cameras and fast interaction. However, this technology is new and installed cameras typically do not incorporate algorithms geared for forensic searches of persons of interest.

The third solution is to perform computations at the edge with an improved person recognition approach (commonly referred to as re-identification, although it is actually re-recognition). A basic color histogram approach does not perform as well as more dedicated systems such as MCHH [9]. Early performance estimates have demonstrated that this system allows human operators to find a person much faster in a database than without the system [9]. For example, as shown in Figure 3, a specialized person search algorithm is much better than a color histogram and five times faster than a manual database search [4]. An operator can find and track a person more efficiently with 37% less misses, which is a significant improvement [8]. Integration of this technology in cameras will significantly enhance the ability of the system to find and track persons. This solution has the advantage that computations can be done at a large scale, enabling the rapid finding of a person using all the cameras. Integration in an installed base is possible if the manufacturer is willing to collaborate on the project or if the cameras have open-source software development kits.

Preparation for this kind of technology can start immediately with its integration in cameras and camera surveillance systems. Testing and validation can be done at an airport during a pilot phase that assesses the added value. The system could be used operationally within two to three years.

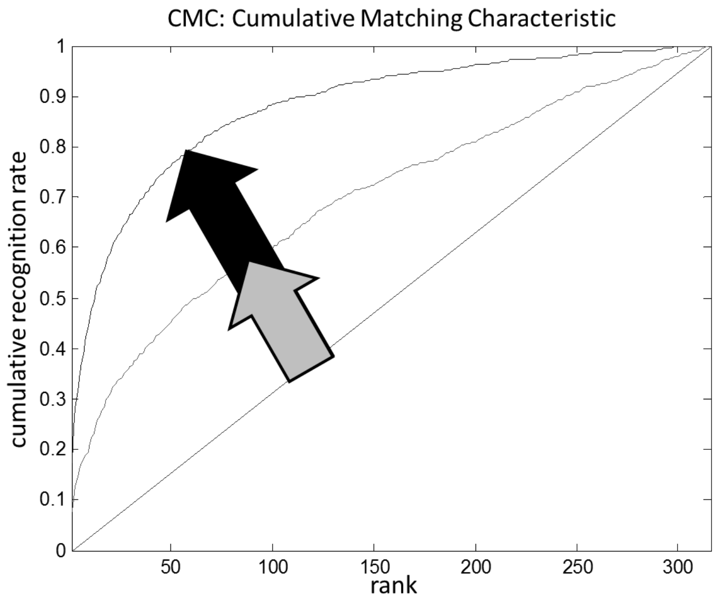

Figure 3: Comparison of person search algorithms [4].

5.4. Failure to capture

All (stand-off) recognition technologies can fail to capture usable data; this may occur by accident or on purpose. Biometric features such as gait are hard to distinguish in crowded environments and reliable face recognition is often not feasible with ordinary CCTV cameras due to their limited resolution. However, even with special high-resolution setups, automated face recognition can be thwarted by individuals who do not look at cameras or wear headgear, coats with high collars or heavy make-up. Recognition based on the complete appearance can be hindered by changing clothes or by wearing clothes that are very similar to those worn by other people. Individuals may actively wish to avoid recognition for reasons of privacy, cultural values or criminal intent. No public empirical data and no verified model are available to predict the actual rates of failure to capture, avoidance and spoofing of recognition by the general public at large airports.

Certain interventions may reduce the rate at which people cannot be recognized. Airport rules could require the faces of individuals to be visible and prohibit clothing that hides facial features. Triggers may be installed that tempt people to look at cameras. Choke points during the departure and arrival processes can be leveraged to capture good images or at least to detect the presence of a person and the failure of a good capture.

6. Information and communications infrastructure

Modern surveillance heavily relies on CCTV cameras and flexible communications depend on mobile networks. An information and communications infrastructure that generates, stores and transports image and data, is, therefore, a key component of any system designed to find and track people in an airport. This section discusses four information and communications technology aspects that help speed up the tasks of finding and tracking people. These include: (i) increased camera coverage; (ii) localization of people; (iii) standardization of surveillance (meta)data; and (iv) sharing information with personnel in the field.

6.1. Increased camera coverage

The finding and tracking of people by operators is largely dependent on camera coverage. Increased camera coverage is constrained by the required financial investment. The first step towards improved coverage is to develop a camera plan that covers specific locations (exits and entrances for all areas, choke points, waiting areas, etc.), networking and data storage, and runtime rules (camera pre-sets and priority user groups).

A key advantage of the widely-used pan-tilt-zoom cameras is that an operator can obtain many details of a large area with limited bandwidth. Disadvantages of pan-tilt-zoom cameras are blind spots in recorded data and camera-use conflicts between multiple users. The disadvantages can be addressed using persistent focused cameras and/or omnidirectional cameras. Initially, these cameras should be preferentially installed at high-risk locations or transition regions (e.g., at entrances and exits) and subsequently at other locations to guarantee minimal coverage at all times in order to improve forensic search. If facial recognition (e.g., mug shot) is used to identify and track a person, then choke points such as entrances and exits should be equipped with high-resolution cameras. Mobile devices can also be used to record close-ups of people. For reliable tracking of people in busy areas (e.g., in queues at customs control), it is advantageous to add cameras that look straight down (e.g., stereo vision systems from Brickstream, Blue Eye Video or Eagle Vision [7]).

6.2. Localization of people

The budgets for security are under constant pressure, leading to more flexibility in observation regions of surveillance personnel and less familiarity with the scenes and camera setup. Intuitive interactions with cameras (Section 4) and multi-camera tracking algorithms (Section 5) could benefit from location and orientation calibrations of cameras and a process for keeping the calibrations up-to-date. This allows mappings of camera pixels to positions in the real-world and vice versa. It is common to perform manual calibrations of fixed cameras. Manual calibrations of pan-tilt-zoom cameras and automatic calibrations of fixed cameras are uncommon.

Another localization aspect that can help improve situational awareness is the indoor localization of security personnel in the field. This can be done using radio-frequency identification, Wi-Fi access points, inertial measurement units, time or angle of arrival, and combinations thereof. Indoor localization enhances the coordination of teams and reduces the time needed to reach a target.

6.3. Standardization of surveillance (meta)data

Large airports are system-of-systems as well as with regard to surveillance systems. Depending on local policies and regulations, sharing CCTV footage or a surveillance infrastructure with multiple organizations may be allowed. For example, airport shops often have their own surveillance systems. When an incident crosses the boundary of a local business and the surrounding airport, relevant images are captured and stored by multiple surveillance systems. Video and metadata standards for surveillance applications such as ONVIF and PSIA have been designed to facilitate interoperability in such situations [41]. Using these interoperability standards also mitigates the risk of vendor lock-in.

The interoperability standards do not support all the required functionality at this time [41]. For example, the template used for finding persons is not standardized, which implies that entire images of people have to be stored, transmitted and processed again by each person search algorithm. A standard for templates would enable the efficient re-use of metadata. Another example is describing the ownership relation between a piece of luggage and its owner. This is not possible in current versions of standards, which hinders the use of new technologies that depend on descriptions of this relation [19].

Airports should implement current metadata schemes in their security systems for operational, entrepreneurial and evaluation purposes. As a community, airports should stimulate the development and use of appropriate metadata schemes by vendors, especially with regard to templates for person search. In addition to the standardization of metadata schemes, it is beneficial for airports to require open software development kits for cameras (as provided by AXIS) and an open surveillance platform that would allow integration of other visual analytics (as provided by Genetec and Milestone). The benefits include the effective use of new technologies and reduced risk due to vendor lock in.

6.4. Sharing information with personnel in the field

Sharing photographs and videos with surveillance personnel in the field greatly improves their effectiveness and efficiency. However, implementing a secure end-to-end distribution platform for pushing photographs and video content to mobile security personnel is a challenge. The required security level of this platform would depend on the classification level set on the transmitted content. Government and industry standards exist for transferring highly sensitive data over wireless links using cryptography (e.g., using the crypto standards Suite B [30, 35] and FIPS 140-2 [29]). Commercial products that meet these standards are available from Aruba networks, Cisco and General Dynamics.

After a secure distribution platform is in place, many applications are available for the exchange of visual content. These applications can enhance interactions between ground personnel and the control room, the use of real-time video, and direct interactions with surveillance cameras and recorded content [27].

Since camera resolution and per-camera coverage are increasing, it is possible to implement functionality for navigating through video content in a spatial manner. Taking the bandwidth limitations of the network and the screen sizes of mobile devices into account, high resolution images need to be tiled for transport and mobile video playback [31]. Subsequently, the application and infrastructure should be developed and audited. The information needs and the possibilities for sharing information should be considered during this process in order to support communications and localization. Initially, this would allow personnel in the field to collaborate more effectively with personnel in the control room. Subsequently, this would enable them to communicate directly with each other using their mobile clients. Figure 4 shows video streams shared with mobile devices independent of their resolution. When a huge high-resolution overview is recorded, the mobile device can focus on a part using limited bandwidth (interested readers are referred to [31] for details). Receiving and sharing visual information should not disturb, but rather assist, personnel in their primary tasks as with simple human-machine interfaces on smart phones, spycams and miniature displays integrated in spectacles (e.g., Google Glass).

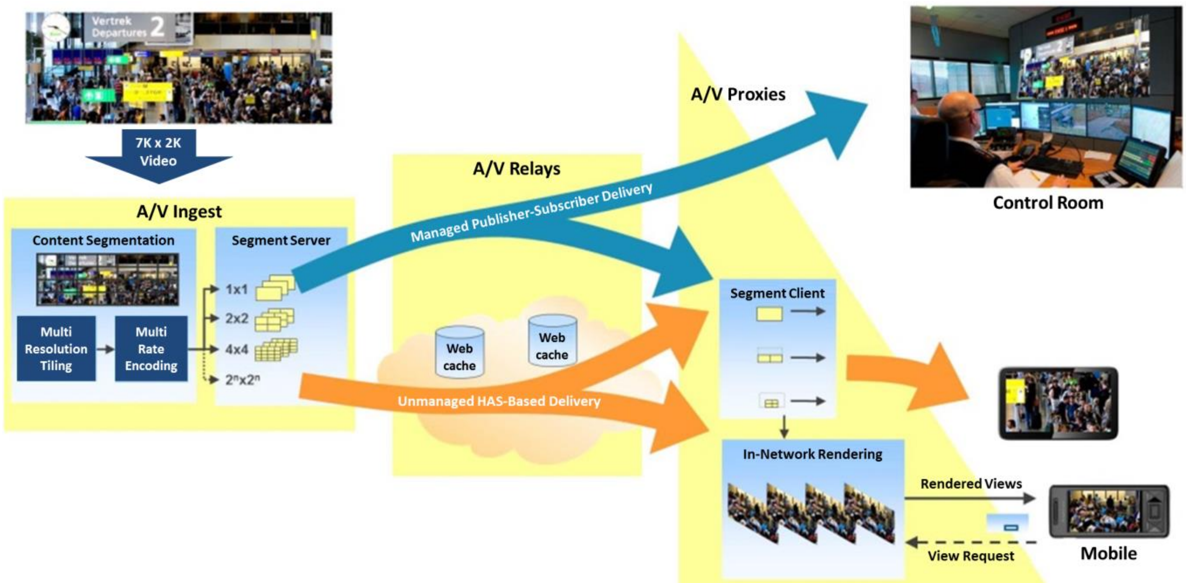

Figure 4: Video streams shared with mobile devices independent of their resolution.

7. Timeline and dependencies in the roadmap

Figure 5 presents the roadmap for the rapid finding and tracking of people at airports. A questionnaire was created to underpin the time estimates in the roadmap. This section describes the questionnaire and summarizes the timeline and dependencies in the roadmap.

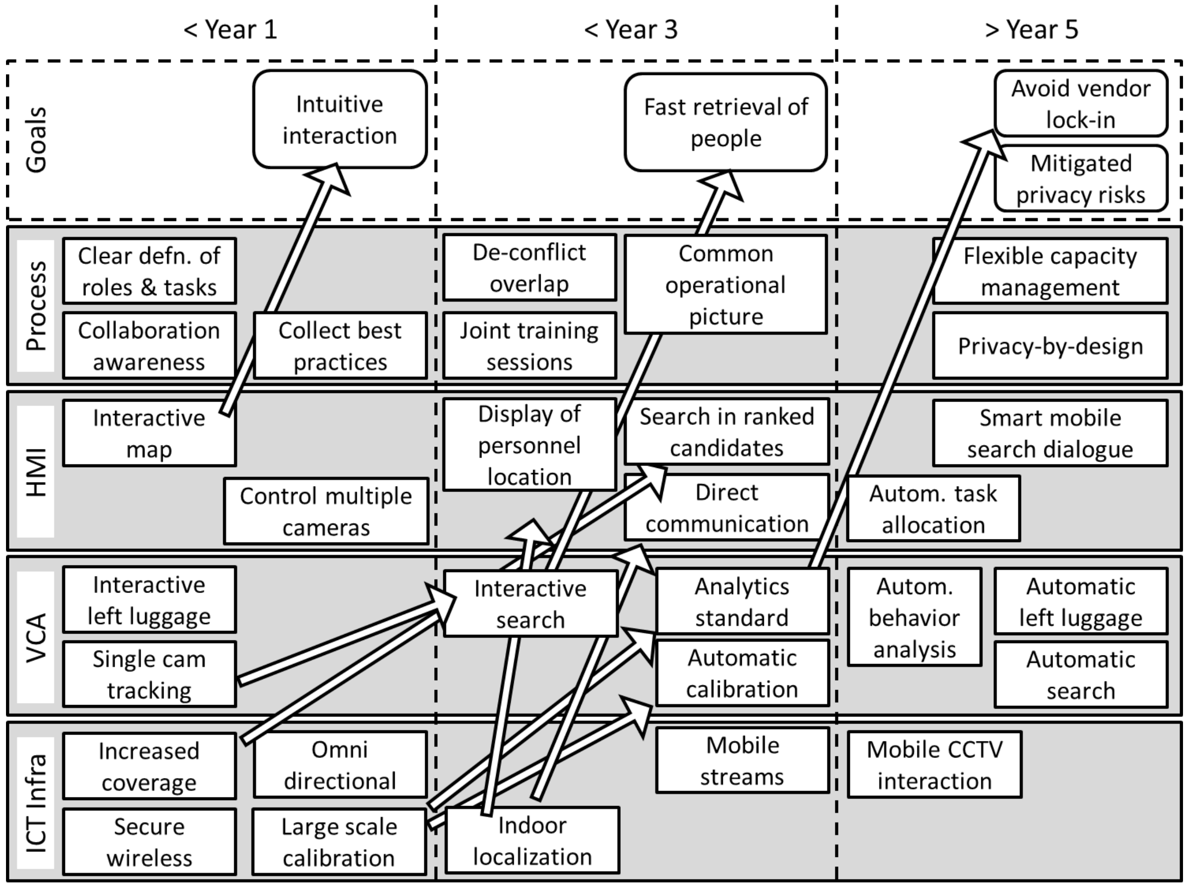

Figure 5: Roadmap for the rapid finding and tracking of people at airports.

7.1. Questionnaire results

A questionnaire was used to underpin the time estimates in the roadmap. Since this paper describes a roadmap for the rapid finding and tracking of people at airports, a large portion of the questionnaire was devoted to these topics. Moreover, emphasis was placed on developments in video content analysis because this is where most of the uncertainty ("hype") exists with respect to performance. The questionnaire also covered techniques that have different levels of maturity in order to support the roadmap timeline. In addition to techniques related to finding and tracking people, several other aspects of video content analysis relevant to the selected use cases at the airport were covered (e.g., detecting left luggage). As described in Section 2, the use cases were selected based on close interactions with stakeholders and end users.

The questionnaire was formulated to obtain quantitative results from eight computer-vision experts for the security and surveillance domain. The respondents were asked whether a capability would be feasible in one year, three years or beyond five years. The questionnaire responses show a distribution over the different time intervals, which is desirable for the roadmap.

Table 1 shows the questionnaire results; note that the totals may not add up to 100 due to rounding. For easy interpretation, the questions are sorted by the maturity indicated by the respondents. Note that the responses to a few questions have wide spreads for the time periods (e.g., Question ID = 4), but many questions have significantly more responses for one of the time periods (e.g., Question ID = 5). High consistency between the respondents and low spreads raise the confidence in and certainty about the estimated maturity of a capability, at least relative to other capabilities.

A binomial test was performed to estimate the confidence intervals. Because of the limited number of samples, it was decided to focus on one category versus the others to draw conclusions; this allows a binomial analysis. Given a binomial test with N=8 samples, the certainty that the majority of responses is correct can be computed. The binomial test shows that, in the case of eight positive responses in a category, it is 99% ($p<0.01$) certain that the majority of experts would vote for the category. Similarly, in the case of zero positive samples, it is 99% certain that the majority would not vote for the category. When there are seven votes or one vote for a category, it is 93% ($p<0.1$) certain about the conclusion. In other cases, it becomes less certain, since for six votes or two votes, the certainty is only 71% ($p>0.1$).

The questionnaire results provide the following maturity rankings in the field of tracking with high certainty (from mature to less mature): (i) automated tracking in a single camera with only one person (100% responses in Year 1 yields 99% certainty that the majority of experts would claim it is feasible in one year); (ii) interactive tracking over multiple cameras without overlapping views where an operator is assisted by a system (88% in Year 3 yields 93% certainty that the majority would vote for this category); (iii) automated finding of a person based on an example image (88% in Year 3 yields 93% certainty) and; (iv) automated tracking of a person in a crowded environment when the person is occluded 50% of the time (100% after Year 5 yields a confidence of 99% that the majority would vote for this category). Thus, it is expected that an interactive semi-automatic system could be used in three years and a full automatic system is not expected to be operationally effective at an airport soon. Note that many of responses depend on the situation in which the technology is used. For example, the capability of fully automatic tracking using a single camera is considered to be very mature (Question ID = 27) as well as very immature (Question ID = 30) depending on the degree to which the environment is crowded.

7.2. Timeline and dependencies in the roadmap

Table 2 summarizes the effects of the developments in the roadmap in terms of the key performance indicators, time and resources. Note that the estimates are very dependent on the uncertainties in time and location. The possibilities within one year are mainly determined by the current starting position of the airport and in three years by what is already technologically feasible.

It is not necessary to develop every component of the roadmap to achieve the desired efficiency. Indeed, there are certain dependencies in the roadmap that should be taken into account. For example, both the computer vision and human-machine interface components rely on information about the positions and orientations of cameras (via the information and communications technology infrastructure) to reduce the search area (computer vision) and to present the results in a geographic context (human-machine interface). Moreover, the incorporation of a novel search technology (computer vision) often requires its integration in a new graphical user interface (human-machine interface) and specialized training on the part of operators to use the system (working process).

Table 1: Questionnaire results based on the responses of eight computer vision experts.

Q-id	Question	% respondents		
		< 1 year	< 3 year	> 5 year
5	Full automatic detection of people passing a virtual tripwire, or people entering or leaving a region.	100	0	0
19	Calibration of fixed staring cameras to allow mapping from image to world coordinates and vice versa.	100	0	0
27	Full automatic tracking of an individual in a single camera when there is (max) 1 person present.	100	0	0
2	Use of left-luggage detection to interactively retrieve the time of placement of a specific piece of luggage that was found by personnel.	88	13	0
3	Full automatic intrusion detection in a sterile zone.	88	13	0
14	Face recognition at special gates that only open after a good face detection.	88	13	0
22	Track people interactively with CCTV cameras with overlapping views.	75	25	0
28	Full automatic tracking of an individual in a single camera when there are (max) 2 persons present.	63	38	0
4	Estimate waiting time in a queue with a high density of stereo cameras on the ceiling that look straight down.	38	38	25
8	Full automatic tracking of an individual in a single camera.	38	38	25
21	Track people automatically with a high density of stereo cameras on the ceiling that look straight down.	38	38	25
11	Automatic pedestrian detection computed on-board inside the camera.	25	75	0
16	Interactively finding a person based on an example image.	25	75	0
20	Calibration of PTZ cameras to allow mapping from image to world coordinates and vice versa.	25	63	13
1	Full automatic left-luggage detection.	25	13	63
10	Interactive tracking of an individual over multiple cameras where an operator is assisted by a system	13	88	0
23	Track people interactively with CCTV cameras without overlapping views.	13	88	0
6	Full automatic detection of loitering.	13	38	50
9	Full automatic tracking of an individual over multiple cameras.	13	25	63
13	Automatic detection of interactions between people based on location, such as 'approach' and 'pass'.	0	63	38
24	Track people automatically with CCTV cameras with overlapping views.	0	63	38
29	Full automatic tracking of an individual in a single camera when it is busy and the person is (max) 10% of the time occluded.	0	63	38
25	Track people automatically with CCTV cameras without overlapping views and a gap between the cameras of 2 meter.	0	50	50
12	Automatic detection of interactions between people based on hand movement, such as 'give' or 'receive'.	0	38	63
15	Face recognition in CCTV cameras from a distance.	0	38	63
17	Interactively finding a person based on a (textual/color) description	0	38	63
26	Track people automatically with CCTV cameras without overlapping views and a gap between the cameras of 20 meter.	0	25	75
18	Full automatically finding a person based on an example image.	0	13	88
7	Full automatic detection of suspicious behavior.	0	0	100
30	Full automatic tracking of an individual in a single camera when it is crowded and the person 50% of the time occluded.	0	0	100

Table 2: Effects of roadmap developments in terms of time and resources.

	Finding	Tracking (live)
Now	Finding a person manually takes a long time	1 person by 2 operators.
1 year	Pilot phase for an interactive search tool.	1 person by 1 operator by using intuitive camera interaction.
3 years	Finding persons twice as fast by using the interactive assisting search tool.	Multiple people by 1 operator with assisted tracking and intuitive interaction.
Later	Practically instantaneous full automatic finding of people based on an example image and/or a description.	Multiple people without an operator (full automatic).

8. Conclusions

Rapid finding and tracking of people at airports is possible, and several new technologies and/or other interventions can make this a reality. At this time, the most promising interventions are improving coverage using persistent or panoramic cameras from the viewpoint of tracking and forensic search, improving multi-camera tracking and interactive forensic searches, connecting separate systems to facilitate intuitive geographical-information-based interactions, creating agreements for information sharing between the various stakeholders, and training people to use the novel technologies and systems.

Predictions for the future cannot be guaranteed and assessments of the maturity of technologies by individuals involved in developing the technologies may lead to funding bias (i.e., a tendency of research conclusions to support the interests of financial sponsors). This is illustrated in the 2013 Gartner Hype Cycle for Emerging Technologies [21], which places content analytics near the top of the peak of inflated expectations, which implies that this technology is currently very overvalued. This research survey polled experts to obtain diverse points of view and reduce stochastic errors. However, there is still no certainty that the assessment is free from systematic error (bias). Although all the estimates may be too optimistic or pessimistic, it is believed that the relative maturities of the technologies have been assessed correctly and that it is possible to state clearly that one type of technology is more mature than another.

Acknowledgements

This research was supported by The Netherlands Top Sector “High Tech Systems and Materials” Roadmap Security under the Passive Sensors Project. The authors also wish to thank the Royal Netherlands Marechaussee and Schiphol Airport Security in Amsterdam for their assistance, and the ERNCIP Thematic Workgroup on Video Analytics and Surveillance for stimulating the idea of using a roadmap for coherent sets of technologies that go beyond a single airport.

References

[1] D. Alberts, R. Huber and J. Moffat, NATO NEC C2 Maturity Model, DoD Command and Control Research Program, Washington, DC, 2010.

[2] R. Aly, R. Arandjelovic, K. Chatfield, and others, The AXES submissions at TRECVID 2013, Proceedings of the TREC Video Retrieval Conference, 2013.

[3] L. An, M. Kafai, S. Yang and B. Bhanu, Reference-based person re-identification, Proceedings of the Tenth IEEE International Conference on Advanced Video and Signal-Based Surveillance, pp. 244–249, 2013.

[4] A. Andriyenko, K. Schindler and S. Roth, Discrete-continuous optimization for multi-target tracking, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1926–1933, 2012.

[5] B. Bechky, Gaffers, gofers and grips: Role-based coordination in temporary organizations, Organization Science, vol. 17(1), pp. 3–21, 2006.

[6] Bosch Security Systems, IVA 5.5 Intelligent Video Analysis, Operation Manual, Grasbrunn, Germany, 2012.

[7] H. Bouma, J. Baan, S. Borsboom, K. van Zon, X. Luo, B. Loke, B. Stoeller, H. van Kuilenburg and J. Dijk, WPSS: Watching people security services, Proceedings of SPIE, vol. 8901(0H), 2013.

[8] H. Bouma, J. Baan, S. Landsmeer, C. Kruszynski, G. van Antwerpen and J. Dijk, Real-time tracking and fast retrieval of persons in multiple surveillance cameras of a shopping mall, Proceedings of SPIE, vol. 8756(0A), 2013.

[9] H. Bouma, S. Borsboom, R. den Hollander, S. Landsmeer and M. Worring, Re-identification of persons in multi-camera surveillance under varying viewpoints and illumination, Proceedings of SPIE, vol. 8359(0Q), 2012.

[10] H. Bouma, G. Burghouts, L. de Penning, P. Hanckmann, J. ten Hove, S. Korzec, M. Kruithof, S. Landsmeer, C. van Leeuwen, S. van den Broek, A. Halma, R. den Hollander and K. Schutte, Recognition and localization of relevant human behavior in videos, Proceedings of SPIE, vol. 8711(0B), 2013.

[11] G. Burghouts, P. Eendebak, H. Bouma and J. ten Hove, Improved action recognition by combining multiple 2D views in the bag-of-words model, Proceedings of the Tenth IEEE International Conference on Advanced Video and Signal-Based Surveillance, pp. 250–255, 2013.

[12] G. Burghouts, K. Schutte, H. Bouma and R. den Hollander, Selection of negative samples and two-stage combination of multiple features for action detection in thousands of videos, Machine Vision and Applications, vol. 25(1), pp. 85–98, 2014.

[13] J. Child and R. McGrath, Organizations unfettered: Organizational form in an information-intensive economy, Academy of Management Journal, vol. 44(6), pp. 1135–1148, 2001.

[14] COPRA Project, Final Results, Fraunhofer Institute for High-Speed Dynamics, Ernst-Mach-Institute, Freiburg, Germany (www. copra-project.eu/Results.html), 2014.

[15] G. DeSanctis and P. Monge, Communication processes for virtual organizations, Organization Science, vol. 10(6), pp. 693–703, 1999.

[16] J. Dijk, Y. Rieter-Barrell, J. van Rest and H. Bouma, Intelligent sensor networks for surveillance, Journal of Police Studies: Technology-Led Policing, vol. 3(20), pp. 109–125, 2011.

[17] M. Endsley, B. Bolte and D. Jones, Designing for Situation Awareness: An Approach to User-Centered Design, Taylor and Francis, London, United Kingdom, 2003.

[18] M. Farenzena, L. Bazzani, A. Perina, V. Murino and M. Cristani, Person re-identification by symmetry-driven accumulation of local features, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2360–2367, 2010.

[19] J. Ferryman, D. Hogg, J. Sochman, A. Behera, J. Rodriguez-Serrano, S. Worgan, L. Li, V. Leung, M. Evans, P. Cornic, S. Herbin, S. Schlenger and M. Dose, Robust abandoned object detection integrating wide area visual surveillance and social context, Pattern Recognition Letters, vol. 34(7), pp. 789–798, 2013.

[20] M. Garcia, The Design and Evaluation of Physical Protection Systems, Butterworth-Heinemann, Burlington, Massachusetts, 2008.

[21] Gartner, Gartner's 2013 Hype Cycle for Emerging Technologies maps out evolving relationship between humans and machines, Press Release, Stamford, Connecticut (www.gartner.com/newsroom/id/2575515), August 19, 2013.

[22] O. Hamdoun, F. Moutarde, B. Stanciulescu and B. Steux, Person re-identification in multi-camera system by signature based on interest point descriptors collected on short video sequences, Proceedings of the Second ACM/IEEE International Conference on Distributed Smart Cameras, 2008.

[23] N. Hu, H. Bouma and M. Worring, Tracking individuals in surveillance video of a high-density crowd, Proceedings of SPIE, vol. 8399(09), 2012.

[24] S. Khan and M. Shah, Tracking multiple occluding people by localizing on multiple scene planes, IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 31(3), pp. 505–519, 2009.

[25] F. Kooi, System for Displaying Surveillance Images, European Patent Application EP 2325820 A1, 2011.

[26] J. Martinez-del-Rincon, J. Herrero-Jaraba, J. Gomez and C. Orrite-Urunuela, Automatic left luggage detection and tracking using multi-camera UKF, Proceedings of the Ninth IEEE International Workshop on Performance Evaluation Tracking Systems, pp. 59–66, 2006.

[27] M. Montagud, F. Boronat, H. Stokking and R. van Brandenburg, Inter-destination multimedia synchronization: Schemes, use cases and standardization, Multimedia Systems, vol. 18(6), pp. 459–482, 2012.

[28] I. Nasrullah, Hierarchical Query Mechanisms for Searchable Encrypted Databases, M.Sc. Thesis, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Delft, The Netherlands, 2009.
[29] National Institute of Standards and Technology, Security Requirements for Cryptographic Modules, FIPS PUB 140-2, Gaithersburg, Maryland, 2001.
[30] National Security Agency, Cryptography Today, Fort Meade, Maryland (www.nsa.gov/ia/programs/suiteb_cryptography/index.shtml), 2015.
[31] O. Niamut, M. Prins, R. Brandenburg and A. Havekes, Spatial tiling and streaming in an immersive media delivery network, Adjunct Proc. European Interactive TV Conference, pp. 84–87, 2011.
[32] R. Nickerson, Confirmation bias: A ubiquitous phenomenon in many guises, Review of General Psychology, vol. 2(2), pp. 175–220, 1998.
[33] G. Roelofsen and B. Gelbord, Method and System for Processing of Image Data, European Patent Application EP 1280117 A1, 2003.
[34] J. Rypkema, M. Neerincx and P. Passenier, Attuning the task design to "envisioned" high-demand situations, Proceedings of the European Safety and Reliability Conference, vol. 2, pp. 1351–1355, 2003.
[35] M. Salter and R. Housley, Suite B Profile for Transport Layer Security (TLS), RFC 6460, 2012.
[36] R. Satta, Dissimilarity-Based People Re-Identification and Search for Intelligent Video Surveillance, Ph.D. Thesis, Department of Electronic and Electrical Engineering, University of Cagliari, Cagliari, Italy, 2013.
[37] R. Satta, G. Fumera and F. Roli, A general method for appearance-based people search based on textual queries, Proceedings of the European Conference on Computer Vision, pp. 453–461, 2012.
[38] Y. Tian, R. Feris, H. Liu, A. Hampapur and M. Sun, Robust detection of abandoned and removed objects in complex surveillance videos, IEEE Transactions on Systems Man and Cybernetics, Part C: Applications and Reviews, vol. 41(5), pp. 565–576, 2011.
[39] W. Treurniet, K. van Buul-Besseling and J. Wolbers, Collaboration awareness – A necessity in crisis response coordination, Proc. Int. Conf. Information Systems for Crisis Response and Management, 2012.
[40] J. van Rest, D. Boonstra, M. Everts, M. van Rijn and R. van Paasen, Designing privacy by design, Proceedings of the First Annual Privacy Forum – Privacy Technologies and Policy, pp. 55–72, 2012.
[41] J. van Rest, F. Grootjen, M. Grootjen, R. Wijn, O. Aarts, M. Roelofs, G. Burghouts, H. Bouma, L. Alic and W. Kraaij, Requirements for multimedia metadata schemes in surveillance applications for security, Multimedia Tools and Applications, vol. 70, pp. 573–598, 2014.
[42] D. Wright and C. Raab, Constructing a surveillance impact assessment, Computer Law and Security Review, vol. 28(6), pp. 613–626, 2012.