VISION-BASED ANALYSIS OF PEDESTRIAN TRAFFIC DATA

Philip Kelly and Noel E. O'Connor

Centre for Digital Video Processing, Adaptive Information Cluster, Dublin City University, Ireland

ABSTRACT

Reducing traffic congestion has become a major issue within urban environments. Traditional approaches, such as increasing road sizes, may prove impossible in certain scenarios, such as city centres, or ineffectual if current predictions of large growth in world traffic volumes hold true. An alternative approach lies with increasing the management efficiency of pre-existing infrastructure and public transport systems through the use of Intelligent Transportation Systems (ITS). In this paper, we focus on the requirement of obtaining robust pedestrian traffic flow data within these areas. We propose the use of a flexible and robust stereo-vision pedestrian detection and tracking approach as a basis for obtaining this information. Given this framework, we propose the use of a pedestrian indexing scheme and a suite of tools, which facilitates the declaration of user-defined pedestrian events or requests for specific statistical traffic flow data. The detection of the required events or the constant flow of statistical information can be incorporated into a variety of ITS solutions for applications in traffic management, public transport systems and urban planning.

1. INTRODUCTION

Traffic congestion has become a considerable problem within urban environments. A classical technique of dealing with this problem involves increasing road capacity by adding more lanes or new routes. However, this approach for expansion, if not properly planned to incorporate future traffic growth – which may be driven by a number of variable factors such as economic growth, urban-sprawl and increasing demand for travel – will only represent a short-term solution to traffic congestion issues. In addition to these issues, if the predictions of an increase in world traffic volumes from 33 trillion passenger kilometres in 2000 to 105 trillion in 2050 [1] is valid, then it may not be economically or environmentally viable to increase road capacity to the required volumes. Furthermore, it may not be possible to address traffic bottlenecks within many cities via these techniques due to the economic cost of purchasing the required land and removing pre-existing, and possibly protected, structures.

An alternative long-term solution to traffic congestion problems may lie in increasing the efficiency of currently existing transportation infrastructures. This can be made possible by; (1) encompassing a broad range of Intelligent Transportation Systems (ITS) [2] into existing infrastructures; and (2), where appropriate, redesigning transportation infrastructures to be better suited to observed traffic-flow characteristics. There is a lot of synergy between these two techniques. The vision of ITS depicts the application of advanced sensor technologies to help monitor and manage traffic flow, improve safety and increase the productivity of transportation systems in real-time. The second technique examines similar traffic data, but retrospectively, with a view to extracting relevant statistical information to help better understand the traffic flow and behavioural information of the area over a given period of time. This information could then be incorporated to help redesign existing infrastructures to increase efficiency, or in the creation of detailed traffic flow models to be used in the design of new infrastructures.

Central to both techniques is the need for a constant inward flow of reliable information describing the amount, type and behaviour of traffic throughout the region to be managed. Traditional approaches for obtaining this data tend to be intrusive. For example, within the area of vehicular traffic monitoring, basic traffic flow information such as counting and classification of the type of vehicles on a road can be provided by locating devices – such as induction loops, passive magnetic sensors or pneumatic tubes [2] – on or under the road. However, with the use of more sophisticated sensor data, more complex applications become plausible, such as redirecting vehicular traffic during adverse conditions or early incident detection that can be used to alert emergency services to a crash as well as inform public road users of dangers ahead. Recently, with the emergence of video cameras as a means to monitor traffic and the development of algorithms to analyse this data, this more sophisticated sensor data is becoming readily available.

To date, research on automated visual-based traffic monitoring systems has generally focused on vehicular traffic. However, vehicular traffic data forms only a part of traffic flow within urban environments. In this paper, we focus on the requirement for the robust detection and analysis of *pedestrian* traffic flow data within a variety of areas. This data is proposed to be acquired via a robust and flexible computer-vision based approach from a single stereo camera. Using this information, the efficiency and safety of pedestrian traffic flow through urban environments could be increased, thus helping to promote walking as a viable, safe and healthy transportation option. In addition, this data could be applied within the area of public transportation infrastructures, to help increase efficiency and reliability. Finally, this information has an application within the design process of urban pedestrianised environments and, if obtained strategically throughout a city, could help within the framework of intelligent city planning of urban walkways, roads and public transportation. As such, this pedestrian traffic flow data could play a strategic part in reducing traffic congestion, both pedestrian and vehicular, therefore leading to increased economic productivity and a reduction of the environmental impact of travel.

The paper is organised as follows: Section 2 gives an overview of the potential advantages of detecting and analysing pedestrian traffic within urban environments. Section 3 details a set of tools and an application independent framework that can be applied to create run-time events within an ITS, or filter results in statistical retrieval applications. Section 4 provides experimental results of the system framework from a number of differing application scenarios. Finally, section 5 details conclusions and future work.

2. PEDESTRIAN DETECTION AND TRACKING

As specific goals of various applications differ, so too does the type of pedestrian data required to be collected and analysed. This data may range from basic traffic data, such as simply the number, average velocity and direction of pedestrians passing through a particular area, to more detailed analysis such as the volume of pedestrian traffic entering, exiting or waiting within predefined areas (such as bus stops or traffic light waiting areas), the peak pedestrian density flow per square metre of pavement, or the classification of pedestrian traffic into the number of adults and children.

Some of this information, if provided in real-time, can be applied to increase the efficiency and safety of urban traffic management control systems. Consider the example of an automated pedestrian traffic light system at a busy traffic intersection. To date, deployed systems tend to be intrusive and inefficient in terms of optimising pedestrian and vehicular traffic flow – a button press is required by a pedestrian to request a lighting change and one pedestrian waiting to cross the road will be given the same treatment as several. However, if the pedestrian traffic numbers waiting to cross can be obtained in real-time, then it could be incorporated with vehicular data into a more efficient urban traffic control management system. This would allow the provision for a better balance between the movement of pedestrianised and motorised traffic through specific intersections and the city as a whole.

Real-time pedestrian traffic flow data can also be applied to the area of public transport systems. Currently, these systems can provide passengers with real-time information regarding public transportation data, such as the time remaining before the next bus to a particular destination. However, the converse is not true, as public transport systems do not have access to real-time passenger data, such as if the numbers waiting at the next bus stop. If this were possible, then a second bus could be summoned from the terminus to cope with overflows. Similarly, taxis lingering at empty ranks could be redirected to those with a number of waiting clients.

Pedestrian data can also be used retrospectively for the statistical analysis of traffic flow patterns through surveillance areas. Pedestrian traffic flow data, including parameters such as density, velocity, direction and delay can be collected over time, analysed and used to make informed decisions about improved infrastructure designs or potential areas requiring expansion. This statistical information is also important for public transportation systems as it can provide the details of the average number of daily, hourly or weekly commuters to specific destinations, allowing an increase in the efficiency in the management of a fleet of vehicles. In addition, using pedestrian traffic density and velocity data, the isolation of bottlenecks within buildings (such as train stations) can be determined and hence the building's structural layout can be redesigned and improved within specific areas deemed to constitute a security hazard. For example, extra exit points can be created in strategic locations in case of a fire or other emergencies. Finally, these pedestrian traffic parameters are required by traffic engineers, architects and town-planners in the planning, design and operation of other facilities such as shopping centres, business areas, airports and pedestrian road crossings [3].

In this paper, we propose a flexible system for the analysis of pedestrian traffic data within an application independent framework. As such, user-defined run-time events can be created and used to provide the data requirements for successful ITS. In addition, via the same framework, retrospective analysis of pedestrian data can be obtained. However, before this framework is introduced a discussion on the requirements and the selection of the underlying pedestrian detection and tracking system is provided for context.

2.1. Pedestrian Detection and Tracking System Requirements

As outlined in section 1, central to any of the proposed applications of pedestrian traffic data is the need for a constant inward flow of reliable sensor information. Traditional techniques for obtaining pedestrian traffic data require a large manual effort. This normally involves employing a group of individuals to count passers by in various key areas throughout the same time period, and using all the separate samples to obtain an estimate of crowd numbers. This is not a very accurate technique, nor is it viable for time critical applications. In addition, important pedestrian traffic data, such as average pedestrian velocity or 3D pedestrian flow patterns can not be obtained via these techniques.

In the past few years, pedestrian detection, counting and tracking has become a popular research topic in the area of computer vision. In addition, stereo-based approaches, such as [4], allow the detection of important 3D data, which can be then applied to reconstruct pedestrian flow and velocity through a 3D scene.

However, obtaining the required pedestrian traffic information via computer vision based techniques poses significant challenges in unconstrained real-world crowded environments. In such scenarios, a multitude of complicating factors has to be taken into account, such as the large variability in pedestrian local and global appearance, occlusion and varying environmental conditions, such as background and lighting changes. We believe that a chosen pedestrian detection algorithm should conform to the following properties;

Flexible The technique should be applicable to a variety of scenarios and camera orientations.

Robust The technique should be able to perform robustly regardless of the environmental conditions, person numbers, orientations or positions.

Scalable Techniques, such as [5], which require a high concentration of cameras to monitor a relatively small spatial region are expensive to install, calibrate, maintain and are generally not considered viable for scalable pedestrian detection of larger surveillance areas.

Economical The economic cost of installing a chosen technique must be less than the savings made by the application of its data. However, it is acknowledged that the cost can be offset by sharing the statistical data obtained from the camera with a variety of third parties, including market research companies.

Real-time If the data is required for time critical ITS, the acquisition of this information may need to be done in real-time.

3D Reconstruction Depending upon the envisioned application, the reliable 3D positioning, velocity or height of pedestrians may need to be obtained.

2.2. Proposed Pedestrian Detection and Tracking System

Many of the pedestrian detection and tracking approaches proposed in the literature produce good results when presented with constrained scenarios that allow specific assumptions to be made. This includes assumptions about the environmental conditions, pedestrian appearance, pedestrian and background colour intensity information, occlusions, or that a person enters the scene un-occluded. Unfortunately, due to these assumptions, few approaches produce reliable results for long periods of time in unconstrained environments [4].

In this work, a robust pedestrian detection and tracking system, originally presented in [6], is applied that augments traditional 2D image processing with stereo vision-based techniques. This single stereo-camera technique specifically targets relatively unconstrained environments and attempts to minimise constraining assumptions. In addition, it requires *no* external training[1] and is robust to; (a) occlusion, even when multiple people enter the scene in a crowd; (b) lack of variability in colour intensity between pedestrians and background; (c) rapidly changing illumination conditions; (d) pedestrians appearing for only a small number of frames; (e) relatively unconstrained pedestrian movement; (f) relatively unconstrained pedestrian pose, appearance and position with respect to the camera; (g) varying camera heights, rotations and orientations; and (h) static pedestrians.

Although the proposed system was designed to minimise constraining assumptions, a small number of inherent assumptions still exist within the system framework. They include that; (1) pedestrians in the scene are standing upright with respect to the groundplane; (2) all moving objects in the scene (within the volume of interest) are caused by foreground pedestrians; and (3) pedestrians in the scene are moving at a velocity of less than 3 metres per second. In addition to this, there are a number of limitations on the type of scenario it can be used in. These include; (a) that a relatively flat groundplane is present within the scene, where no object of interest is located below this groundplane; (b) the camera must be orientated so that the groundplane is visible in the image plane; and (c) the system is only able to reliably detect pedestrians for a short-medium range, up to a maximum distance of 8 metres from the camera. Finally, the proposed system is not currently real-time, with the processing time of a single 640×480 pixel frame taking between 10–20 seconds on a standard 2GHz laptop. Obviously this is far from real-time processing and, as such, all experiments are currently implemented using pre-recorded data-set sequences. However, in future work a number of optimisation techniques are envisioned that will significantly reduce the algorithmic complexity of the proposed system.

Evaluation of this system on 10,000 ground-truth pedestrians from a number of sequences with varying camera height, camera orientation and environmental conditions reveal an accurate performance of the proposed approach (94.1% precision and 84.6% recall) [6]. In addition, the system was evaluated using 3D methodologies that resulted in an average error of 8.94cm in the 3D positioning of detected pedestrians and an average error in 3D height of 10.02 cm (this equates to an average of 2.50% and 5.60% error in the average positioning and height respectively) [6].

Figure 1(a)-(d) presents some illustrative results from a variety of scenarios where each detected pedestrian is enclosed by a bounding box of a certain colour. An example of a tracking sequence can also be seen in figure 1 (e)-(h), where

[1]However, it is acknowledged that the designer has brought in his own area of expertise into setting a number of hard-coded thresholds throughout the system framework.

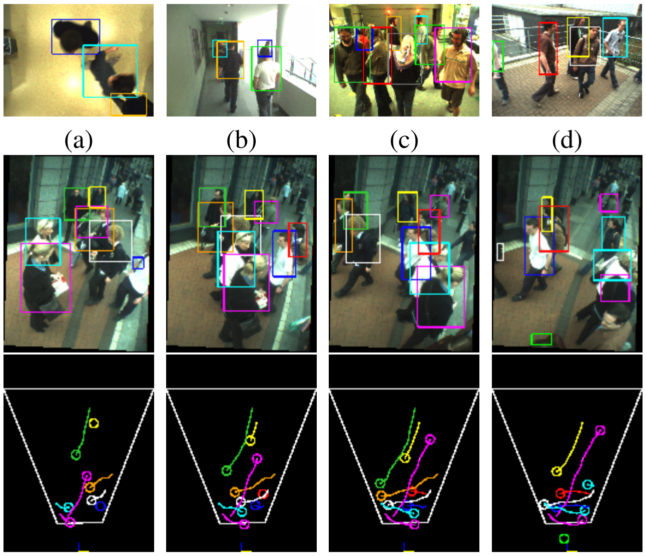

(a) (b) (c) (d)

(e) (f) (g) (h)

Fig. 1. Examples; (a)-(d) Pedestrian detection from the *Overhead*, *Corridor*, *Vicon* and *DCU Corner* scenarios respectively; (e)-(h) Pedestrian tracking from the *Grafton* scenario.

the second row of images depict the scene from a *plan-view* or *birds-eye view* orientation. In these plan-view images, the white lines indicate the bounds of the scene (that are defined with respect to the visible groundplane within the scene), the position of detected pedestrians in that frame are illustrated by a circle of the same colour as their bounding box and pedestrian tracks – i.e. where they have previously been within the scene – are depicted as "tails" from the centre of the circle to their temporal scene positions.

3. PROPOSED SYSTEM FRAMEWORK

The pedestrian detection and tracking technique, outlined in section 2, can provide rich pedestrian traffic data required by a number of applications. As the proposed technique also obtains a 3D reconstruction of pedestrian position, the statistical output provided could also incorporate more detailed analysis, such as average pedestrian velocity, height or density flow in user-defined 3D directions. However, for many applications scenarios a further level of processing is required. For example, within the automated pedestrian traffic light system outlined in section 2, not all pedestrian traffic in the scene is relevant. Clearly detecting pedestrians is a necessary preprocessing step, but just because a person is in the scene does not mean that they want to cross the road. However, if the person walks towards the crossroads, stops and waits, then this can probably be assumed to be the case.

As such, a specific application requirement may be to examine pedestrian traffic for specific scenarios, or *events*, such as when the number of pedestrians within a certain area of the scene exceeds a certain threshold. This type of event detection may be required either in a real-time ITS application *or* in a more sophisticated statistical retrieval and analysis application (for example, the user may require to know the number of pedestrians waiting at a crossroads over a specific time period). However, for many flexible applications the exact event detector *cannot* be hard-coded into the system framework as either; (1) the event definition is dependent on an undefined scene; or (2) the event is itself undefined.

The first scenario is typical of many Ambient Intelligence (AmI) [7] applications, such as the automated pedestrian traffic light system, as although the event required to be detected is known (i.e. detect static pedestrians in a designated area waiting to cross the road), information about the scene (i.e. the exact designated crossroads area) is unknown and may vary depending upon the camera positioning and scene specific properties. The second scenario is typical of the envisioned statistical pedestrian traffic retrieval application whereby a user may require the statistical data to be filtered by a specific constraining event. However, the details of this event may remain unspecified until a specific point is reached during the execution of the application (i.e. until the user declares the required event).

Our framework allows a variety of user-defined applications and events to be created and tailored by an end-user during run-time (i.e. during the execution of the system framework) for use in ITS. In addition, the same framework can be applied to statistical retrieval applications, whereby results can be filtered by user-defined events during run-time. Within this framework, a pedestrian *indexing* scheme is applied whereby a surveillance video is augmented with robust pedestrian tracking and statistical information – such as 3D position and height data – using the technique outlined in section 2. For real-time event detection applications, such as the automated pedestrian traffic light system, a list of all pedestrians currently in the scene, plus their tracking and statistical information is maintained within the system framework. From this information, specific events can be detected. In addition, for statistical retrieval applications the full augmented video can be quickly searched retrospectively for specific events via the augmented meta-data.

In order to create tailored events for use within a variety of application scenarios, an event syntax is incorporated into the system framework. This syntax provides a suite of event detectors that can be declared, parsed, and run during run-time. These events are associated with a concept called a *hotspot*.

3.1. Hotspots

3D hotspot regions are created from 2D plan-view images – such as that depicted in figure 2(a), which illustrates the surveillance area of the *DCU Corner* scenario of figure 1(d) from a *plan-view* orientation. In figure 2(a), it should be noted that a background colour model has been projected onto the image-plane to allow the gauging of distance and orientation within the plan-view image. Using this plan-view image, an arbitrarily shaped hotspot region is simply created in the proposed system framework by allowing the user to draw a region of interest within the plan-view image – an example of a hotspot is depicted as a yellow coloured area in figure 2(f). The resultant hotspot, extended downwards towards the groundplane, can be seen as a 3D area of interest.

Using these hotspots, events and statistical output can be constrained to incorporate specific types of pedestrian tracks – for example, tracks that start-on, start-off, or pass-through a hotspot, or pedestrians that pass through the hotspot within a narrow range of directions, such as only those travelling in a northerly direction. Any number of hotspots can be created and incorporated into a specific application. If more than one is created, then logical operators can then be applied between them – thus, for example, it becomes possible to obtain all those pedestrians who pass-through multiple hotspots – or conversely, obtain all the tracks that do not. These hotspot regions can be used for tailoring specific applications to differing scenarios. For example, within the automated pedestrian traffic light application, the area where pedestrians are expected to wait can be quickly and easily defined.

3.2. Event Syntax

The framework syntax provides a suite of event detectors that can be applied to analyse pedestrian traffic data. Currently, the syntax incorporates a single type of core-event that triggers when a pedestrian adheres to the underlying hotspot's properties for the previous N seconds (e.g. if the hotspot is set so that only pedestrians currently-on the hotspot are detected, then for each frame the associated hotspot event will trigger once for each pedestrian who has remained on the hotpot for the previous N seconds). Although this single core-event is relatively simple, it provides the means within the framework to create and tailor a suite of user-defined pedestrian traffic monitoring events and applications. This is possible as within framework syntax the value of N and a variety of hotspot properties can be set – including global logical associations between hotspots, and hotspot filters based on direction, position (e.g. if a pedestrian track must start-on the hotspot, etc) or velocity, which can be applied to obtain specific types of pedestrian tracks. For each change in the underlying hotspot's parameters (and N) a new event definition can be declared.

4. EXPERIMENTAL RESULTS

In this section, it is demonstrated how the framework can be tailored to both event-driven ITS and statistical retrieval applications.

4.1. Statistical Analysis

The proposed system framework can be applied to statistical retrieval applications. In such a scenario, pedestrian traffic data from a specific period of time is examined retrospectively for the purpose of studying its behaviour over time. In this section, two specific examples of how the system framework can be applied within this area are presented; (a) the first obtains the most popular entrance and exit points within a scene; and (b) the second examines pedestrian traffic flow paths in and out of specific areas within a scene.

4.1.1. Pedestrian Traffic Entry Points

An example of the use of hotspots is illustrated in figure 2. In this example, figure 2(b) depicts 191 independent pedestrian tracks from a 3 minute test sequence (for an example of the camera position in this scene, see the *DCU Corner* scenario of figure 1(d)). On their own, these pedestrian tracks provide little useful information about pedestrian traffic flow within the scene. Little more can be obtained from an image of the average pedestrian footfall in the scene – see figure 2(c) – where the brighter the red colour, the more pedestrians have traversed that scene position. However, figures 2(d) and (e) – which illustrate each pedestrian track starting point in green – provide slightly more informative data. From this data, it can be seen that the starting points of pedestrians entering the scene form a "V" at the bottom of the image (the other green starting points tend to be caused by fully occluded, and temporarily lost, pedestrians).

An illustration of the use of hotspots to obtain useful statistical information from this data can be seen in figures 2(f)-(j). In figure 2(f), a hotspot is created around the area where most pedestrians enter the scene from the right-hand side. Using the system framework syntax, the properties of the hotspot is set so that only pedestrians that start-on this hotspot are incorporated into the retrieved pedestrian track results – see figures 2(g) and (h), which depict the resultant tracks that start-on this hotspot and the average pedestrian footfall from these tracks respectively. In addition, the results can be filtered so that the footfall image incorporates only the parts of those remaining tracks that appear on a hotspot – see figure 2(i). From this resultant image, the brighter the red colour, the more often that position has been traversed by entering pedestrians. By recursively filtering this image (achieved by recursively squaring the average pedestrian footfall at each point) the most popular entering point of pedestrians becomes clear – see figure 2(j). This statistic is interesting when a similar hotspot on the opposing side is similarly examined – see figures 2(k)-(o). By comparing figures 2(j) and (o), the most popular entering points of pedestrians from both sides can be compared. From these figures it can be seen that on the right-hand side there is a single most popular entrance point into the scene, however on the left-hand side there are two similarly popular entering points located close together but lower down

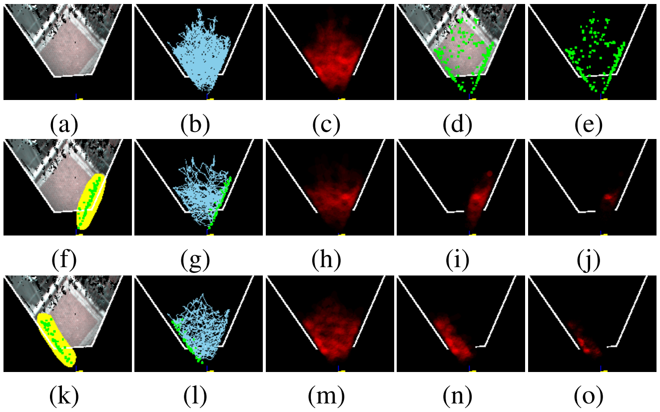

Fig. 2. Hotspot Statistics Usage; (a) Plan-view image; (b)/(g)/(l) Pedestrian tracks; (c)/(h)-(j)/(m)-(o) Average pedestrian footfall; (d)/(e) Pedestrian track starting points; (f)/(k) Hotspots.

in the plan-view image (i.e. closer to the camera). Clearly, the entrance points between the two sides of the scene are asymmetric.

4.1.2. Pedestrian Traffic Flow Lanes

A second example of the use of hotspots in the retrieval of statistical information is illustrated in figure 3. In this example, 201 independent pedestrian tracks from a three 1 minute test sequences are used. Unlike the previous scenario, the three sequences from this scene – referred to as the *Grafton* scenario, see figures 1(e)-(h) – were taken from a camera that, with the help of Dublin County Council, was mounted on a traffic light pole on Grafton Street, a busy pedestrianised shopping street in Dublin city centre. The sequences taken from this scenario are of real-world data pedestrian traffic data, of relatively high traffic flow, from the general public walking during their daily routine.

Figure 3(a) illustrates the scene from a plan-view orientation with a background colour model projected onto the image-plane. To orientate the reader, pedestrians tend to enter/exit the scene from three positions; (1) the first – see figure 3(b) – is from a distance, where pedestrians approach the camera parallel to its principal axis; (2) the second – see figure 3(c) – is from the left, where pedestrians are coming from/going towards Grafton Street; and (3) the third – see figure 3(d) – from the right, where traffic has to wait at a pedestrian crossing controlled by a traffic light system.

As with the previous sequence, little information can be obtained from all pedestrian tracks or the total average footfall in the scene (see figures 3(e) and (f) respectively). However, using hotspots very interesting information about pedestrian flow paths can be extracted from the scene. In figure 3(row1)(g) a hotspot is created that covers the first area where pedestrians enter the scene. If this hotspot is set so that only pedestrian tracks that start-on and end-off the hotspot are included in the statistics, the tracks and average pedestrian traffic footfall displayed in figures 3(row1)(h) and (i) are respectively obtained. However, if the same hotspot is set so that only pedestrian tracks that start-off and end-on the hotspot are included, the tracks and average traffic footfall of conforming pedestrians differ significantly – see figures 3(row1)(j) and (k). Note that the average footfall in figure 3(k) are coloured green for illustrative purposes only. By overlaying the two average footfalls – as in figure 3(row1)(l) – the two average traffic footfalls can be directly compared. In this figure, two distinct pedestrian traffic flow lanes can be seen, one for pedestrians entering the hotspot area (depicted in red) and a second for pedestrians exiting the same hotspot area (depicted in green).

Figure 3(row1)(l) illustrates a self-organization phenomena within pedestrian traffic, described in [8], where at sufficiently high densities, pedestrians form lanes of uniform walking direction. In addition, this figure also conforms earlier empirical observations of pedestrian streams [9], which cites that pedestrians develop an asymmetric avoidance behaviour with respect to the right or left-hand side. In [9] it was determined that in Germany, pedestrians tended to the right-hand side to avoid on-coming traffic. From figure 3(row1)(l) it can be seen that in Ireland, pedestrians also tend towards the right-hand side.

Further asymmetric pedestrian traffic flow lane pairs can be found if similar analysis is made upon the other two areas where pedestrians enter the scene, see figures 3(row2-row3)(g)-(l). Note in all these images, the red and green areas represent pedestrians entering and exiting the relevant hotspot respectively, as such the colour of similar flow lanes between figures 3(row1-row3) may change, depending upon the relevant hotspot. In each case, it is found that pedestrians tend towards the right-hand side when avoiding on-coming traffic. It is noted however, that analysis of the third exit/entrance area within the scene – see figure 3(row3)(l) – differs from the previous two scenarios in that there is significantly more green than red. This scenario is caused by there being significantly more pedestrians walking out of the hotspot area compared to those entering it. This result is to be expected within the context of the scene, as most people exiting this area are walking up towards the busy pedestrianised shopping street (i.e. walking from exit 3 to exit 2), whereas people travelling the other direction can either walk towards exits 1 or 3 and depending upon their end destination they may choose either. As such, far fewer pedestrians travel from exit 2 to exit 3 with respect those who travel in the opposite direction, leading to a higher saturation of green areas in figure 3(row3)(l).

4.2. Event Detection

As outlined in section 3.2, a hotspot can be used in conjunction with a system framework syntax which allows user-

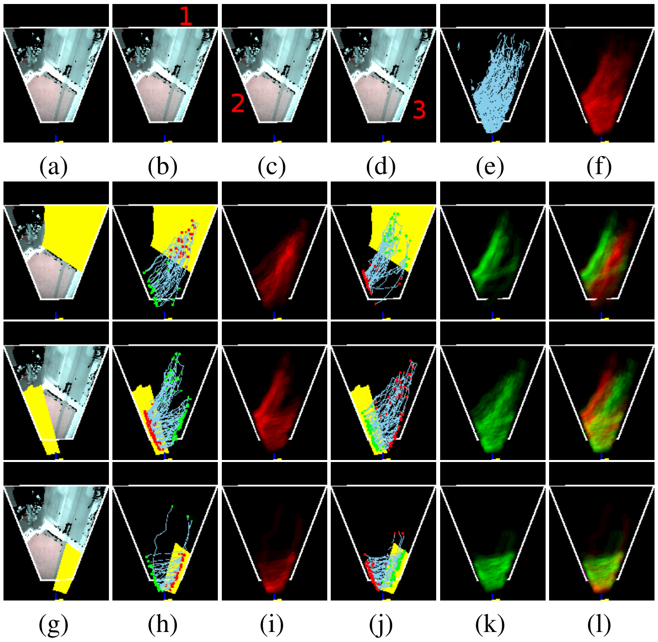

Fig. 3. Pedestrian flow paths; (a) Plan-view image; (b)-(d) Entrances/Exits 1-3; (e) Tracks; (f) Average footfall; (g) Hotspot; (h) Tracks starting-off and ending-on hotspot; (i) Average footfall from (h); (j) Tracks starting-on and ending-off hotspot; (k) Average footfall from (j); (l) Average footfall from (h) and (j).

defined events to be created. In this section, two specific event types are created by varying the core-event and underlying hotspot parameters; (a) the first event can be used to create entrance counts in user-defined areas; and (b) the second determines the number of pedestrians standing for a given period of time within a certain area.

4.2.1. Pedestrian Traffic Entrance Counts

The first event declaration scenario involves the use of single user-defined hotspot region of interest. Within the framework two separate events are created (both with the core-event parameter N set to 1); (1) the first triggers when a pedestrian first enters the hotspot (i.e. when a pedestrian was outside the hotspot in the previous frame and is inside in the current frame); and (2) the second triggers when a pedestrian first exits the same hotspot. In each scenario, a single count is held within the system framework for the number of times a given event triggers. Using these two simple event definitions, traffic counts entering and exiting specific areas of interest can be easily obtained. This type of event can be used in conjunction with the previous analysis of pedestrian flow at scene entry points, or it can be applied in real-time to keep a running count of pedestrian numbers currently occupying a building.

An example of this application scenario is presented in figure 4. It should be noted, that although the camera orientation within this sequence differs significantly to others in this work, the same underlying pedestrian detection and tracking algorithms have been applied, i.e. the system does not have to be trained/tailored for each specific camera orientation. In figures 4(e)-(h) the detected pedestrians in row 1 are surrounded by a red bounding box if they have not triggered an event, and a blue or green bounding box if they have entered or exited a hotspot respectively, thus altering the total entrance counts for the scene hotspot. In a similar manner to figure 1(e)-(h), the plan-view tracks of the detected pedestrians are depicted in figures 4(e)-(h) row 2.

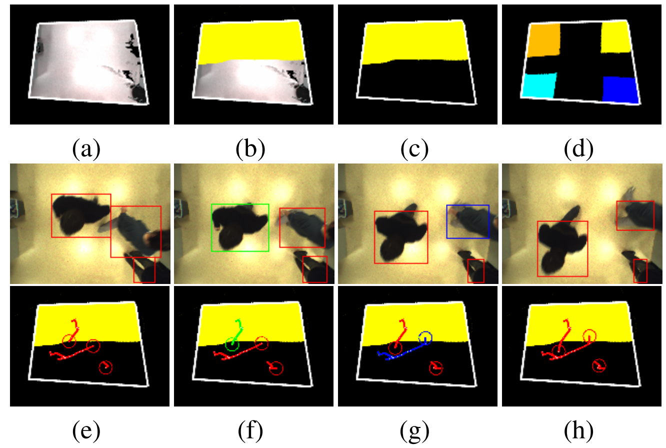

Fig. 4. Entrance Event; (a) Plan-view image with background; (b) Hotspot with background; (c) Hotspot; (d) Multiple hotspots; (e)/(h) No triggered events; (f) Pedestrian exiting hotspot; (g) Pedestrian entering hotspot.

Finally, it should be noted that although this example had one hotspot, a number of them can be used in co-operation. For example, in figure 4(d) four separate hotspots are used and logically OR-ed together. As such, the counts are based on pedestrians entering or exiting any of the hotspot regions. Using this technique, this application can be simply extended to monitor a scene with multiple sparsely located hotspots.

4.2.2. Pedestrian Traffic Area Waiting Counts

The second event declaration scenario involves the use of a single hotspot, where the event triggers when a pedestrian stands-on the underlying hotspot for a given period of time (defined by the core-event parameter N) within the hotspot area. This has a variety of applications, such as determining the number of people waiting in all areas of public transport, such as bus stops, taxi ranks and train stations. In addition it can be used to detect lingering pedestrians in surveillance video. However, the demonstration application in this section is that of the automated pedestrian traffic light system described in section 2. In this application, an event is triggered if a pedestrian enters the scene, walks towards the area to cross, stops and waits for a pre-defined period of time. In

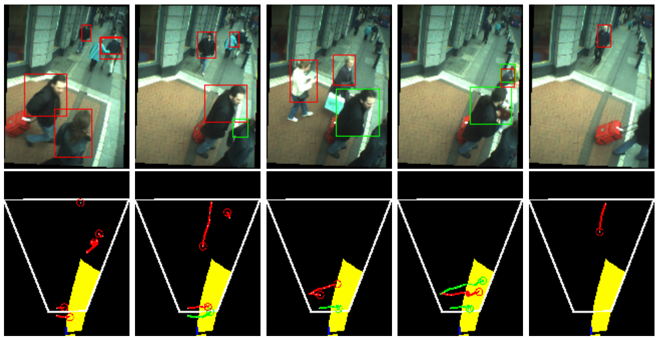

(a) (b) (c) (d) (e)

Fig. 5. Waiting Event; (a)/(e) No triggered events; (b)/(c) 1 pedestrian waiting; (d) 2 pedestrians waiting.

our experiments, we defined N to be 1.5 seconds and that pedestrians are defined to be stopped if their velocity drops below 0.3 metres per second (this was achieved within the event by setting a hotspot filter to ignore all pedestrians with a velocity of greater than 0.3 metres per second). This velocity value was chosen as it would allow a pedestrian's 3D position to fluctuate by up to 0.15 metres from its correct position in any given frame – as outlined in section 2.2 this is nearly twice the expected error of the proposed system.

An example of this application scenario is presented in figure 5 where a hotspot is drawn, similar to that of figure 3(Row 3)(g), via the 2D plan-view image. As before, in figures 5(e)-(h) the pedestrians are depicted in green if they have triggered an event (i.e. are waiting to cross the road), otherwise they are shown in red. In this sequence, two pedestrians enter from the left (see figure 5(a)) and travel to the crossroads hotspot waiting area where they stop. Initially the closer of the two triggers an event in figure 5(b), but that person then moves slightly forwards and out of range of the stereo-cameras overlapping field of view. As such, the pedestrian no longer gets detected and its track is lost. However, the second person then triggers an event after standing still for the allotted time period – see figure 5(c). This pedestrian remains still for a period of time until two more enter the hotspot region; the farthest from the camera then stops and triggers and event, while the second keeps walking and does not – see figure 5(d). Eventually the traffic lights change and the pedestrians cross the road in figure 5(e).

5. CONCLUSIONS AND FUTURE WORK

In this work, we presented a flexible system for the analysis of pedestrian traffic data within an application independent framework. From this system framework a variety user-defined run-time events can be created and used to provide the data requirements for successful ITS and statistical retrieval applications. Currently, only a single core-event definition is incorporated into the framework syntax – in future work the syntax will be extended to incorporate a wider variety of core-event definitions. In addition, further features such as time of day, pedestrian interactions (such as pedestrians walking in a group or on their own) and pedestrian statistics (such as colour and height) will be incorporated into event definitions. In addition, a hierarchical hotspot clustering framework will be developed so that more complex logical operators can be applied between hotspot regions, and thus more sophisticated events can be declared.

Acknowledgements

This material is based on works supported by Science Foundation Ireland under Grant No. 03/IN.3/I361.

6. REFERENCES

[1] A. Schäfer, "Long-term trends in global passenger mobility," *U.S. National Academy of Engineering*, vol. 36, no. 4, 2006.

[2] K.A Lawrence, *Sensor Technologies and Data Requirements for ITS (Intelligent Transportation Systems)*, Artech Houe, 2001.

[3] A. Rourke and M.G.H. Bell, "An image-processing system for pedestrian data collection," in *International Conference on Road Traffic Monitoring and Control*, 1994, pp. 123–126.

[4] M. Harville, "Stereo person tracking with adaptive plan-view templates of height and occupancy statistics," *International Journal of Computer Vision*, vol. 22, pp. 127–142, 2004.

[5] A. Mittal and L.S. Davis, "M2tracker: A multi-view approach to segmenting and tracking people in a cluttered scene using region-based stereo," in *European Conference on Computer Vision*, 2002, vol. 1, pp. 18–36.

[6] P. Kelly, *Pedestrian Detection and Tracking using Stereo Vision Techniques*, Ph.D. thesis, Dublin City University (DCU), URL http://elm.eeng.dcu.ie/˜kellyp/thesis, 2008.

[7] P. Remagnino and G.L. Foresti, "Ambient intelligence: A new multidisciplinary paradigm," in *IEEE Transactions on Systems, Man and Cybernetics*, 2005, vol. 35, pp. 1–6.

[8] D. Helbing, P. Molnár, I.J. Farkas, and K. Bolay, "Self-organizing pedestrian movement," *Environment and Planning B: Planning and Design*, vol. 28, no. 3, pp. 361–383, 2001.

[9] S.J. Older, "Movement of pedestrians on footways in shopping streets," *Traffic Engineering and Control*, vol. 10, pp. 160–163, 1968.