

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 April 2008 (24.04.2008)

PCT

(10) International Publication Number
WO 2008/047077 A1

(51) International Patent Classification:

G06T 7/20 (2006.01) **G06K 9/00** (2006.01)
G06K 9/64 (2006.01) **G06K 9/62** (2006.01)
G08B 13/194 (2006.01)

(21) International Application Number:

PCT/GB2007/003843

(22) International Filing Date: 10 October 2007 (10.10.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

0620353.3 16 October 2006 (16.10.2006) GB
06270091.9 16 October 2006 (16.10.2006) EP

(71) Applicant (for all designated States except US): **BAE SYSTEMS PLC** [GB/GB]; 6 Carlton Gardens, London, SW1Y 5AD (GB).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **WILLIAMS, Mark, Lawrence** [GB/GB]; BAE Systems Plc, Atc Sowerby, PO Box 5, Filton, Bristol BS34 7QW (GB).


(74) Agent: **BAE SYSTEMS PLC**; Group IP Department, P.O. Box 87, Farnborough Aerospace Centre, Farnborough, Hampshire GU14 6YU (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL,

[Continued on next page]

(54) Title: IMPROVEMENTS RELATING TO EVENT DETECTION



(57) Abstract: A method of event detection for detecting the occurrence of an event in a crowd is disclosed. The method includes the steps of providing video footage of the crowd, which video footage includes a series of image frames; calculating a series of values representative of the crowd dynamics, which series of values is calculated in dependence on the series of image frames; analysing the series of values to identify any discrepancies occurring in the series of values; and, in response to identifying a discrepancy, providing an output signal to alert an operator that an event has been detected. In preferred embodiments, the series of values are related to the complexity of the crowd dynamics. Apparatus suitable for carrying out such a method is also disclosed.



PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report*

- 1 -

IMPROVEMENTS RELATING TO EVENT DETECTION

This invention relates to a method and apparatus for detecting the occurrence of events in crowds, which events are captured in video footage (such as closed circuit television (CCTV) footage). The present invention may be used either to monitor video footage in real time, and alert an operator when an event occurs, or may be used to analyse pre-recorded video footage in order to search for events.

There are a great many 'events' that merit detection by security systems aimed at counter-terror and anti-social behaviour. One important subset is that of events which occur in or near crowds, since crowds are often either the target, or reason, for events of interest. There may be a wide variety of diverse events of interest that it would be desirable to detect. An area of particular interest and concern is the problem of detecting unforeseen events, since, a priori, it may not be known even what type of event it is desirable to detect. Events that are understood, or expected, can be detected using a set of heuristic rules, but this is not possible where the type of event to be detected is unknown, and, moreover, becomes impractical when there is a desire to detect many different types of event. Events of interest may include terrorist attacks; fires, earthquakes, or other disasters; the beginnings of social unrest; the presence of unusual packages; or even one person acting unusually.

There are a number of well-known complexity measures that relate to the effective number of degrees of freedom in a system, or the dimensionality of a data set. One known complexity measure is described in by W. A. Wright, R. E. Smith, M. Danek, and P. Greenway in their paper "A Generalisable Measure of Self-Organisation and Emergence", published at pages 857 to 864 of "Artificial Neural Networks - ICANN 2001" by Springer-Verlag (2001) and edited by G. Dorffner, H. Bischof, and K. Hornik. Wright et al. associate a system's ability to exhibit emergent behaviour with sudden transitions in their complexity measure, relative to smooth changes in system parameters. In this context, emergent behaviour may include, for example, the sudden dispersion of a

- 2 -

school of fish in response to the arrival of a predator. Wright et al. suggest that their measure could be used to evaluate and encourage desirable behaviours in a variety of complex systems, including artificial neural systems, evolutionary computation systems, and reinforcement learning systems. However, they
5 unfortunately fail to realise how their complexity measure could be applied in other fields.

It is an object of the present invention to provide an improved method and apparatus for event detection. It is a further object of the present invention to provide a method and system for event detection that overcomes, or at least
10 mitigates, at least some of the above mentioned problems associated with prior-known event detection systems and methods.

In broad terms, the present invention resides in the realisation that analysis of crowd dynamics can be applied to the detection of events occurring in crowds. This enables an event of interest to be detected through the reaction
15 of a crowd to that event, rather than by direct observation of the event. Event detection systems and methods in accordance with the present invention are therefore sensitive to a wide variety of events: the invention provides an event detector that is sensitive to the way in which the crowd reacts, as reflected through changes in the crowd motion, rather than a specific type of event.

20 Therefore, in accordance with a first aspect of the present invention, there is provided a method of detecting the occurrence of an event in a crowd, the method comprising the steps of:

- (a) providing video footage of the crowd, the video footage comprising a series of image frames;
- 25 (b) calculating, for each of the series of image frames, a value representative of the crowd dynamics, the value being calculated in dependence on said each of the series of image frames, thereby calculating a series of values;
- (c) analysing the series of values to identify any discrepancies
30 occurring in the series of values; and

- 3 -

(d) in response to identifying a discrepancy, providing an output signal to alert an operator that an event has been detected.

Thus the reactions of the crowd to an event, analysed in terms of the crowd dynamics, are used to detect the occurrence of that event. The crowd
5 has numerous proximal points of view with which to monitor its surroundings, and in addition the crowd might react to smells, such as smoke, and to audible cues, such as warnings or explosions. Such smells and audible cues cannot be directly detected by a video surveillance system. Moreover, the crowd can identify peculiar or unusual behaviour using subtle clues such as body language
10 that cannot be reliably read from a more remote viewing. The 1995 nerve gas attack on the Tokyo underground is an example of an event – the release of a toxic gas - that could have been detected indirectly through the reaction of those persons exposed to the gas, but not directly through video footage. By analysing the dynamics of crowd, and monitoring the dynamics in real time
15 through use of some value representative of the dynamics, the invention allows a wide variety of events to be detected as discrepancies in the representative value.

The step of analysing the series of values may comprises the step of providing a normal range of values, and a discrepancy may then be identified
20 when the value is outside the normal range. The normal range may be a predetermined range, or may be determined in dependence on the mean and standard deviation of the series of values. Advantageously, this simplifies the process of identifying discrepancies, and is readily implemented automatically by appropriate processing means such as a computer.

25 In one embodiment, the method further comprises the steps of tracking a number of targets moving within the crowd to determine a number of target tracks, and updating the target tracks in response to each successive image frame; wherein the step of calculating a series of values comprises determining a value in dependence on the updated target tracks. The updated target tracks
30 are a direct measure of the crowd dynamics, and therefore little useful information is lost in the processing of the series of images. Conveniently, the step of calculating a series of values comprises the step of forming a track

- 4 -

matrix comprising rows, each row representing a target track, and evaluating Wright's complexity metric for the track matrix. Herein, the term 'Wright's complexity metric' is used to refer to the complexity metric derived by Wright *et al.* in their above-referenced paper.

5 Alternatively, the step of calculating a series of values may comprise calculating a series of difference images in dependence on the series of image frames, the difference image for a current image frame being related to differences between the current image frame and a preceding image frame. Conveniently, the preceding image frame may be the image frame immediately
10 preceding the first image frame in the series of image frames. Calculation of the series of values representative of the crowd dynamics from a series of difference images, rather than from a series of target tracks, avoids the problem of requiring complex tracking techniques to track individual targets within the crowd, thus simplifying the event detection method.

15 Conveniently, the step of calculating the series of values in dependence on the difference images may comprise, for each of the series of values, the steps of: forming a difference matrix, the difference matrix comprising a number of rows and each row representing one of the series of difference images; and evaluating Wright's complexity metric for the difference matrix to provide said
20 each of the series of values. The method may further comprise the step of lowering the resolution of at least one of the series of image frames prior to calculating the series of values, advantageously enhancing the efficiency of the processing of the difference matrix.

Optionally, the method is performed in real time.

25 Conveniently, the series of values comprises a series of complexity values. Complexity values can be related to the number of effective degrees of freedom exhibited in the crowd dynamics. The complexity value is a scalar quantity that can be easily monitored.

30 The invention extends to a computer program product comprising computer readable instructions, which product is operable, when inserted into a

- 5 -

computer in operable association with a video camera, to perform a method as claimed in any preceding claim.

According to a second aspect of the invention, there is provided an event detection system comprising a video camera, a processor, and an alert means; the video camera being operable to record video footage, and to transmit said video footage to the processor, said video footage comprising a series of image frames; and the processor comprising means to calculate, for each of the series of image frames, a value representative of the crowd dynamics, the value being calculated in dependence on said each of the series of image frames, such that a series of values is calculated; and means to analyse the series of values to identify any discrepancies occurring in the series of complexity values; and wherein the alert means is configured to alert an operator in response to the occurrence of any discrepancies identified by the process

The above and further features of the invention are set forth in the appended claims, and will be explained in the following, by way of example only, with reference to various specific embodiments which are illustrated in the accompanying drawings, in which:

Figure 1 is a schematic view of apparatus according to a first embodiment of the invention;

Figure 2 is a flow diagram schematically illustrating steps performed in operation of the embodiment of Figure 1;

Figure 3 is a graph illustrating data recorded during the test of the method illustrated in the flow diagram of Figure 2; and

Figure 4 is a flow diagram schematically illustrating steps performed as part of a method in accordance with a further embodiment of the invention.

Referring firstly to Figure 1, there is shown apparatus 100 in accordance with a first embodiment of the invention. Apparatus 100 comprises a video camera 110, processor 120, and alert means 130. Video camera 110 is arranged to monitor an area of interest in which there is a crowd. The area of interest may be a crowded high street, an airport terminal, or even a road on

- 6 -

which traffic is flowing. Thus it will be understood that the term 'crowd', as used herein, should be construed broadly to include any collection of targets of interest, and also to include collections of targets of interest that move through an area, whether or not the targets are simultaneously present in that area.

- 5 Camera 110 may be any type of device operable to record a series of images of its field of view, as will be understood by those skilled in the art, but in this particular embodiment, camera 110 is a CCTV camera.

Video footage, in the form of a series of image frames, is output from camera 110 to processor 120. In the present embodiment, processor 120
10 monitors the images it receives in real time in order to detect the occurrence of events. For each received image, a complexity value is calculated, so that the processor 120 calculates a series of complexity values. The complexity is a useful quantity to measure because it is representative of the crowd dynamics. More particularly, it is related to the effective number of degrees of freedom
15 present in the movement of the members of the crowd. Consider a simple situation in which system 110 is monitoring a crowd of people in a large hall. Under normal circumstances, people in the crowd will move largely at random, with little correlation between the paths taken by individuals. However, should there be, for example, a fire at one end of the room, people in the crowd would
20 begin to move towards exits, and there would be much greater correlation between the movement of individuals within the crowd. In the latter situation, the effective number of degrees of freedom of movement of the crowd is lowered, and thus the fire would be detected as an event. Notably, the event would be detected even if the fire were not in the field of view of camera 110.
25 The manner in which the calculation of the complexity value is performed will be described in detail hereinafter.

The standard deviation and mean of the series of complexity values are calculated, and a normal range is defined as two standard deviations either side of the mean. If the value calculated for a given image frame is outside this
30 normal range, processor 120 provides a signal, via alert means 130, to an operator. In this specific embodiment, alert means 130 is an audible alarm. As explained above, deviation from the normal range of complexity values can be

- 7 -

associated with the occurrence of an event of interest, and the alert provided to the operator allows the operator to investigate more closely. For example, where the operator must monitor the output of not just camera 110, but also several cameras, the alert may cause the operator to look at the footage
5 recorded by camera 110 more closely.

The operation of processor 120 is shown schematically in flow diagram 200, which illustrates the steps carried out by processor 120 in determining whether an alert should be given in response to the receipt of a new image frame. Flow diagram 200 therefore takes, as its starting point, the receipt of a
10 new image frame by processor 120, which step is illustrated as step 210. In the following, this newly received image frame is referred to as the current image frame. At step 220, a difference image is calculated. The difference image is indicative of how the current image frame has changed relative to a preceding frame. Images can be represented in the form of a matrix of the pixel values,
15 and the difference image is calculated by subtracting, from the matrix representing the current image frame, the matrix representing the image frame received immediately prior to the current image frame.

The calculation of a difference image at step 220 considerably simplifies subsequent computations to be carried out on the images, since pixels
20 unchanged between the two frames will all, in the difference image, be represented by a matrix element having a value of zero. Moreover, this removes redundant, background information from the current image frame. Processing can be further simplified by the application of a threshold to the difference image matrix, such that elements having a value below a preset
25 threshold value are given a value of zero, and elements having a value above a preset threshold value are given a value of one. The difference image matrix is thus mapped to a binary matrix, and consists of elements having a value of either 1 or 0. In some circumstances, and particularly where only few pixels change values between frames, it is advantageous to interchange the elements,
30 such that elements previously having a value of 0 are changed to have a value of 1, and vice versa. This is because there can be numerical instabilities in the complexity calculation described below (and in particular at the step of singular

- 8 -

value decomposition) where the matrix has a large population of zero-valued elements.

At step 230, a complexity value is calculated using data from a number of difference images stored by processor 100, and data from the difference image
5 calculated from the current image frame. The stored difference images are those immediately preceding the current image frame. Processor 120 stores a constant number of difference images so that, in response to the current image frame being received (at step 210) and its associated difference image being
10 calculated (at step 220), the oldest difference image stored is deleted. In the present embodiment, twenty difference images are stored at any one time. The complexity value is then calculated from a difference matrix Z defined by:

$$Z = \begin{pmatrix} \overset{1}{x_{t=t_0}^1} & \overset{2}{x_{t=t_0}^2} & \overset{3}{x_{t=t_0}^3} & \cdots & \overset{N}{x_{t=t_0}^N} \\ \overset{1}{x_{t=t_0+1}^1} & \overset{2}{x_{t=t_0+1}^2} & \overset{3}{x_{t=t_0+1}^3} & \cdots & \overset{N}{x_{t=t_0+1}^N} \\ \vdots & & & & \\ \overset{1}{x_{t=T}^1} & \overset{2}{x_{t=T}^2} & \overset{3}{x_{t=T}^3} & \cdots & \overset{N}{x_{t=T}^N} \end{pmatrix}$$

in which the element $x_{t=j}^n$ is the value of the n th pixel in the difference image recorded at $t = j$. Thus, the matrix representing the difference image at $t = j$ is
15 converted into a row of pixel values, and the difference matrix Z is comprised of nineteen rows representing difference images calculated from a series of nineteen consecutive image frames received immediately prior to the current image frame, and the difference image calculated for the current image frame. The difference image calculated from the current image frame is that labelled
20 $t = T$, and the oldest difference image is that labelled $t = t_0$.

In effect, therefore, as each new difference image is calculated, a complexity value is evaluated using a window of image frames, which window is of constant size $(T - t_0)$ and, in this specific embodiment, is equivalent to twenty image frames. As will be clear to those skilled in the art, the time index
25 is measured in terms of the number of image frames received. As will also be clear to those skilled in the art, it is necessary to initialise difference matrix Z when monitoring begins, but how this is done is not critical since the

- 9 -

initialisation procedure does not affect the calculations once more than $(T - t_0)$ images have been recorded by video camera 110. One possible initialisation technique that is used in the present embodiment is to define an initial matrix Z_0 , of an appropriate size, whose elements are all valued at unity.

5 The complexity value is evaluated using Wright's complexity metric:

$$\Omega = 2^S$$

where

$$S = \sum_{i=1}^M \sigma'_i \log_2 \sigma'_i$$

and σ'_i is the i th normalised singular value of difference matrix Z . The
 10 normalised singular values are obtained using singular value decomposition, a well-known operation applicable to any matrix that yields a diagonal matrix having diagonal entries that are the singular values of the original matrix. For any real matrix A , the singular values are the square roots of the eigenvalues of the matrix A^*A , where A^* denotes the transpose of matrix A . These values are
 15 normalised using a simple scaling procedure that scales the largest singular value of matrix Z to 1. The complexity metric Ω is derived in the paper by Wright *et al.* referred to above, and can be viewed as representing the effective number of degrees of freedom exhibited by the crowd.

The calculation of Ω , and particularly the step of singular value
 20 decomposition leading to the values σ'_i , can be numerically lengthy where the number of pixels in each image frame is large. The required processing time for this calculation can be reduced by lowering the resolution of the image frames, thus reducing the number of pixels. Since each element in the difference matrix Z is related to a pixel in a difference image, the number of elements in each row
 25 of the difference matrix Z , and therefore the overall size of difference matrix Z , is reduced by lowering the resolution of the input image frames. This shortens the amount of processor time required for the calculation of Ω , largely because of the reduction in the amount of time needed for the step of singular value

- 10 -

decomposition. It can be determined by trial and error how much the resolution of the images can be reduced, but it is noted that the competing requirements of speed of processing, and sufficiency of resolution to capture information relating to the crowd dynamics, must both be satisfied.

5 At step 240, the complexity evaluated at step 230 is compared to a normal range. The normal range, as described hereinabove, is defined to be within predetermined threshold values, set by an operator, that vary according to the standard deviation of the series of complexity values. In this particular embodiment, the normal range is chosen to be a range of two standard
10 deviations either side of the mean of the complexity values. If the complexity value evaluated at step 230 is within the normal range, processor 120 returns to step 210 and awaits a further new image frame. If the complexity value is outside the normal range, processor 120, at step 250, provides an output to alert means 130 such that an operator is alerted to the discrepancy in the series
15 of complexity values. The operator can then investigate further the causes of the discrepancy, for example by focusing attention on the output from camera 110.

 The method illustrated by flow diagram 200 shown in Figure 2 has been demonstrated to work successfully using sample data from a car park. An
20 artificial event was created by placing a "scarecrow" in the car part approximately half-way through the trial period. Figure 3 is a graphical representation of the variation of the complexity measure with time, as measured using a method in accordance with the second embodiment of the invention. The presence of the scarecrow alters the path that persons walking
25 across the car park take, and this is seen in graph 300 of Figure 3 as a transition (labelled 310) in the variation of the complexity with time.

 The transition 310 is clearly a statistically significant transition of the type that would be readily detectable by automatic means. In fact, the shift in the mean complexity value on placing the scarecrow in the car park is
30 approximately equivalent to thirteen standard deviations. The complexity measure was evaluated, in this example, using an image resolution of approximately one pixel per person, equivalent to approximately 500 pixels per

- 11 -

frame. Such a resolution was found to be sufficiently computationally efficient and still allowed for satisfactory event detection. Lowering the resolution further was found to have the effect of lowering the significance of the transition 310 in the complexity measure, since the standard deviation of the series of complexity
5 values increased. Thus the detection of any discrepancies in the series of values was rendered less likely by lowering the resolution of the image frames.

As described above, the complexity measure derived by Wright et al, is related to the dynamics of the objects to which it is applied. It is therefore surprising that the application of Wright's complexity metric to a matrix Z
10 comprising elements representing pixel values should result in any meaningful value. However, it is thought that the difference images do, in fact, retain some information on the crowd dynamics which can be extracted through application of the complexity metric. Information contained within the difference images relates only to those pixels whose values have changed between consecutive
15 frames. These changed pixels are clearly indicative of movement within the crowd, since any such movement will cause pixel values to change. Thus the difference matrix Z may be successfully used in combination with the complexity measure for the purposes of event detection.

Furthermore, whilst the 'event' in the car park example above, of a
20 scarecrow appearing next to a parked car, is entirely artificial, those skilled in the art will immediately recognise that there are many circumstances in which a crowd of people would react similarly to an event worthy of detection. For example, a crowd would be likely to react to the presence of a suspicious unattended package: individuals in the crowd would be likely to alter their paths
25 in order to move further away from the package. Such a reaction would be detected by a system according to the first embodiment of the invention. Notably, it is the reaction of the crowd, rather than the presence of the suspicious package, that is detected. A system able to automatically detect such events has clear utility in the field of homeland security.

30 A method in accordance with a further embodiment of the invention is illustrated schematically by flow diagram 400 shown in Figure 4. The method of this embodiment uses the apparatus 100 illustrated in Figure 1, and is similar to

- 12 -

the method of this embodiment illustrated schematically by flow diagram 200 shown in Figure 2, except in the processing of the image data used to calculate the complexity measure. In the flow diagram 400 of Figure 4, step 410 represents a new image frame being received by processor 120 (Figure 1). The
5 image is taken to be of an area in which there are a number of targets of interest moving, which targets may, for example, be cars moving in traffic on a busy street. The targets are tracked by processor 120, and their tracks recorded and updated in response to the receipt of each new image frame. The step of updating the target tracks is represented as step 420 in flow diagram
10 400. The targets may be tracked using any known tracking algorithm, such as a Kalman Filter algorithm (disclosed by R. E. Kalman in "A New approach to linear filtering and prediction problems", published in the Transactions of the ASME Journal of Engineering, pages 33 to 45 (1960)), or a particle filter algorithm (disclosed by N. Gordon, D. Salmond, and A. Smith in their paper
15 "Novel Approach to non-linear/non-Gaussian Bayesian State Estimation", published in IEE Proc. F, volume 140, no.2, pages 107 to 113 (1993)), in conjunction with a suitable target detection method, such as, for example, a template-matching technique. The template-matching technique requires a processor to attempt to match a template learned in advance to a part of a
20 received image. Where a match is found, a target is identified.

At step 430, the complexity of the tracks is evaluated using Wright's complexity measure Ω defined above. In this case, the matrix Z comprises rows of data that represent the target tracks rather than data from the difference images, and is therefore referred to in the following as a track matrix. Thus, for
25 example, each row may comprise coordinates, in the image frame, of each target, with the time evolution of the tracks being recorded as new rows added to the track matrix with each new image received. The complexity value is then calculated as above, using Wright's complexity metric. It is noted that lowering the resolution of the images before calculating the complexity value is
30 unnecessary in this case, since there is no link between image resolution and the size of the track matrix. In fact, it is expected that a reduction in image resolution would have the undesirable effect of lowering the reliability of the

- 13 -

tracking algorithm used to track targets through the video footage. At step 440, it is determined whether or not the complexity value lies within a defined normal range, as described above in relation to step 240 of flow diagram 200 illustrated in Figure 1. If the complexity value is within the normal range, the process
5 returns to step 410. If it is outside of the normal range, an operator is alerted to the discrepancy at step 450, and the process returns to step 410.

The method of the embodiment illustrated in Figure 4 is also applicable to targets other than cars, and, in particular, is applicable to crowds of people. The method can therefore be quantitatively compared to that illustrated in
10 Figure 2, and the inventors have shown that the Figure 4 method can produce better results than the Figure 2 method. A comparative test, analysing the same video data shown in Figure 2 using the method of the embodiment illustrated in Figure 4 showed a shift of 22, rather than 13, standard deviations of the complexity value on placing the scarecrow in the car park. This is
15 thought to be because the track data is more closely related to the dynamics of the target crowd than the pixel-based data. However, in this example, the tracking of people moving across the car park was performed manually, rather than using any known tracking algorithms, because known tracking algorithms are frequently unreliable for tracking people. People can change shape as they
20 move, can change their appearance significantly by turning and changing the aspect presented to the camera, and can be partially or completely obscured by the objects in the frame. Thus it can be very difficult to successfully track a person through a series of image forms because of problems associated with target detection. Template matching techniques (as described briefly above)
25 can be used, as can model-based database techniques, in which use prior knowledge of the known skeletal system, and the manner in which it moves, to impose constraint on target detection. However, such techniques can require significant processing power.

A method in accordance with a yet further embodiment of the invention is
30 illustrated schematically in flow diagram 500 shown in Figure 5. The method can be performed using apparatus 100 shown in Figure 1 where camera 110 is operable to output colour image frames to processor 120. The steps illustrated

- 14 -

in flow diagram 500 are similar to the steps illustrated in flow diagram 200 shown in Figure 2, except in that the image is colour quantised at step 520 before difference images and complexity values are calculated. Thus, at step 510, a new image frame is received at processor 120 (shown in Figure 1). At
5 step 520, the image frame is colour quantised. In the colour image frame, each pixel has a colour value. The colour quantisation process assigns a value of 1 or 0 to each pixel in dependence on whether its colour lies between a predetermined range of colour values. A number of colour quantised image matrices can therefore be produced, covering the spectrum of colours to which
10 camera 110 is sensitive. Difference images can then be calculated, at step 530, for each of the colour quantised images, and thus a difference matrix can be calculated for each colour range. Thus each new image frame generates a number of difference matrices, complexity values for which are evaluated at step 540 using Wright's complexity metric defined above. The various
15 complexity values for each of the difference matrices are summed in order to arrive at a single complexity value for the image frame. At step 550, it is determined whether or not the single complexity value is within a normal range (in the same way as is done at step 240 in the method illustrated in flow diagram 200). If the value is within the normal range, the processor returns to
20 step 510 and awaits a new image frame; if the complexity value is without the normal range, an operator is alerted at step 560. The inclusion of the step of colour quantisation has been found to result in a more effective event detection method in some trials using real footage.

Having described the invention with reference to various specific
25 embodiments, it is noted that these embodiments are in all respects exemplary. Variations and modifications are possible without departing from the scope of the invention, which is defined in the accompanying claims. Such variations and modifications will be immediately obvious to those skilled in the art. For example, whilst, in the above, it has been described to define a normal range in
30 terms of the standard deviation and mean of the series of complexity values, it would of course also be possible for an operator to provide predetermined

- 15 -

thresholds, or indeed for a discrepancy to be identified when the value changes by more than a predefined amount between two (or more) image frames.

It has also been described in the above to calculate a series of complexity values in dependence on the series of image frames using Wright's complexity metric Ω , but it will be immediately recognised by those skilled in the art that other complexity measures, or, indeed, any measure representative of the crowd dynamics could be used in place of Wright's complexity measure Ω . Such values may be other values related to the number of effective degrees of freedom, or may be marginal statistical values. Examples include Lyapunov exponents, the mean velocity of the targets of interest making up the crowd, or the length of the average straight-line path of the targets of interest making up the crowd. Moreover, it is to be understood that the invention can be applied both in real time, as described above, in order to alert an operator to the occurrence of an event, and in later analysis of video footage captured at an earlier time. The ability to analyse previously obtained video footage is useful for footage acquired from security cameras, since it can significantly reduce the amount of manual labour required to identify a sequence of interest in the footage.

Finally, it is noted that it is to be clearly understood that any feature described above in relation to any one embodiment may be used alone, or in combination with other features described, and may also be used in combination with one or more features of any other of the embodiments, or any combination of any other of the embodiments.

- 16 -

CLAIMS

1. A method of detecting the occurrence of an event in a crowd, the method comprising the steps of:
 - (a) providing video footage of the crowd, the video footage comprising a series of image frames;
 - (b) calculating, for each of the series of image frames, a value representative of the crowd dynamics, the value being calculated in dependence on said each of the series of image frames, thereby calculating a series of values;
 - (c) analysing the series of values to identify any discrepancies occurring in the series of values; and
 - (d) in response to identifying a discrepancy, providing an output signal to alert an operator that an event has been detected.
2. A method as claimed in claim 1 wherein the step of analysing the series of values comprises the step of providing a normal range of values, and wherein a discrepancy is identified when the value is outside the normal range.
3. A method as claimed in claim 2 wherein the range is determined in dependence on the mean and standard deviation of the series of values.
4. A method as claimed in any preceding claim further comprising the steps of tracking a number of targets moving within the crowd to determine a number of target tracks, and updating the target tracks in response to each successive image frame; and wherein the step of calculating a series of values comprises determining a value in dependence on the updated target tracks.
5. A method as claimed in claim 4 wherein the step of calculating a series of values comprises the step of forming a track matrix comprising rows, each row representing a target track, and evaluating Wright's complexity metric for the track matrix.

- 17 -

6. A method as claimed in any of claims 1 to 3 wherein the step of calculating a series of values comprises calculating a series of difference images in dependence on the series of image frames, the difference image for a current image frame being related to differences between the current image frame and a preceding image frame.
7. A method as claimed in claim 6, wherein the preceding image frame is the image frame immediately preceding the current image frame in the series of image frames.
8. A method as claimed in claim 6 or claim 7 wherein the step of calculating the series of values comprises, for each of the series of values, the steps of:
- forming a difference matrix, the difference matrix comprising a number of rows and each row representing one of the series of difference images; and
- evaluating Wright's complexity metric for the difference matrix to provide said each of the series of values.
9. A method as claimed in any one of claims 6 to 8, further comprising the step of lowering the resolution of at least one of the series of image frames prior to calculating the series of values.
10. A method as claimed in any preceding claim wherein the method is performed in real time.
11. A method as claimed in any preceding claim wherein the series of values comprises a series of complexity values.
12. A method of event detection substantially as described herein with reference to the accompanying drawings.
13. A computer program product comprising computer readable instructions, which product is operable, when inserted into a computer in operable association with a video camera, to perform a method as claimed in any preceding claim.

- 18 -

14. An event detection system comprising a video camera, a processor, and an alert means; the video camera being operable to record video footage, and to transmit said video footage to the processor, said video footage comprising a series of image frames; and the processor
- 5 comprising means to calculate, for each of the series of image frames, a value representative of the crowd dynamics, the value being calculated in dependence on said each of the series of image frames, such that a series of values is calculated; and means to analyse the series of values to identify any discrepancies occurring in the series of complexity values;
- 10 and wherein the alert means is configured to alert an operator in response to the occurrence of any discrepancies identified by the processor.
15. An event detection system as claimed in claim 13 configured to operate in real time.
- 15 16. An event detection system substantially as described herein and with reference to the accompanying figures.

1/5




FIGURE 1

2/5




FIGURE 2

3/5




FIGURE 3

4/5




FIGURE 4




FIGURE 5

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2007/003843

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06T7/20 G06K9/64 G08B13/194 G06K9/00 G06K9/62

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06T G06K G08B

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	BOGHOSSIAN B A ET AL: "Motion-based machine vision techniques for the management of large crowds" ELECTRONICS, CIRCUITS AND SYSTEMS, 1999. PROCEEDINGS OF ICECS '99. THE 6TH IEEE INTERNATIONAL CONFERENCE ON PAFOS, CYPRUS 5-8 SEPT. 1999, PISCATAWAY, NJ, USA, IEEE, US, vol. 2, 5 September 1999 (1999-09-05), pages 961-964, XP010361627 ISBN: 0-7803-5682-9 Section 1. Introduction Section 3. Detection and Segmentation of Crowd Flow Section 4. Detection of Crowd-Related Emergencies ----- -/--	1,2,6, 12-16



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

Z document member of the same patent family

Date of the actual completion of the international search

24 January 2008

Date of mailing of the international search report

05/02/2008

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Klemencic, Ales

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2007/003843

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WRIGHT W.A., SMITH R.E., DANEK M., GREENWAY P.: "A Generalisable Measure of Self-Organisation and Emergence" LECTURE NOTES IN COMPUTER SCIENCE: ARTIFICIAL NEURAL NETWORKS - ICANN 2001, vol. 2130, 21 August 2001 (2001-08-21), - 25 August 2001 (2001-08-25) pages 857-864, XP002424497 cited in the application abstract Section 1. Introduction Section 3. Entropic Measure	1-16
A	DAHMANE M ET AL: "Real-Time Video Surveillance with Self-Organizing Maps" COMPUTER AND ROBOT VISION, 2005. PROCEEDINGS. THE 2ND CANADIAN CONFERENCE ON VICTORIA, BC, CANADA 09-11 MAY 2005, PISCATAWAY, NJ, USA, IEEE, 9 May 2005 (2005-05-09), pages 136-143, XP010809019 ISBN: 0-7695-2319-6 abstract Section 4. Behavioral recognition	1-16
A	HU W ET AL: "A Survey on Visual Surveillance of Object Motion and Behaviors" IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS: PART C: APPLICATIONS AND REVIEWS, IEEE SERVICE CENTER, PISCATAWAY, NJ, US, vol. 34, no. 3, August 2004 (2004-08), pages 334-352, XP011114887 ISSN: 1094-6977 abstract page 334, left-hand column, line 1 - right-hand column, line 5 page 334, right-hand column, line 27 - line 46 Section II.B Motion Segmentation page 338, left-hand column, line 8 - line 21 Section IV. Understanding and Description of Behaviours Section VII. F Anomaly Detection and Behaviour Prediction ----- -/--	1,2,4,6, 12-16

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2007/003843

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	VELASTIN ET AL: "A motion-based image processing system for detecting potentially dangerous situations in underground railway stations" TRANSPORTATION RESEARCH. PART C, EMERGING TECHNOLOGIES, PERGAMON, NEW YORK, NY, GB, vol. 14, no. 2, April 2006 (2006-04), pages 96-113, XP005584782 ISSN: 0968-090X abstract page 98, line 3 - line 30 page 100, line 17 - page 103, line 8 page 106, line 1 - page 107, last line page 110, line 9 - page 111, line 12 -----	1,12-16