

Activity based surveillance video content modelling

Tao Xiang*, Shaogang Gong

Department of Computer Science, Queen Mary, University of London, London E1 4NS, UK

Received 4 April 2006; received in revised form 11 September 2007; accepted 26 November 2007

Abstract

This paper tackles the problem of surveillance video content modelling. Given a set of surveillance videos, the aims of our work are twofold: firstly a continuous video is segmented according to the activities captured in the video; secondly a model is constructed for the video content, based on which an unseen activity pattern can be recognised and any unusual activities can be detected. To segment a video based on activity, we propose a semantically meaningful video content representation method and two segmentation algorithms, one being offline offering high accuracy in segmentation, and the other being online enabling real-time performance. Our video content representation method is based on automatically detected visual events (i.e. ‘what is happening in the scene’). This is in contrast to most previous approaches which represent video content at the signal level using image features such as colour, motion and texture. Our segmentation algorithms are based on detecting breakpoints on a high-dimensional video content trajectory. This differs from most previous approaches which are based on shot change detection and shot grouping. Having segmented continuous surveillance videos based on activity, the activity patterns contained in the video segments are grouped into activity classes and a composite video content model is constructed which is capable of generalising from a small training set to accommodate variations in unseen activity patterns. A run-time accumulative unusual activity measure is introduced to detect unusual behaviour while usual activity patterns are recognised based on an online likelihood ratio test (LRT) method. This ensures robust and reliable activity recognition and unusual activity detection at the shortest possible time once sufficient visual evidence has become available. Comparative experiments have been carried out using over 10h of challenging outdoor surveillance video footages to evaluate the proposed segmentation algorithms and modelling approach.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Video content analysis; Activity recognition; Surveillance video segmentation; Dynamic Bayesian networks; Dynamic scene modelling; Unusual activity detection

1. Introduction

The rapid increase in the amount of CCTV surveillance video data generated has led to an urgent demand for automated analysis of video content. Content analysis for surveillance videos is more challenging than that for broadcasting videos such as news and sports programmes because the latter are more constrained, well structured and of better quality. This paper aims to address a number of key issues of surveillance video content analysis:

1. How to construct a representation of video content which is informative, concise and able to bridge the gap between the

low level visual features embedded in the video data and the high level semantic concepts used by human to describe the video content.

2. How to segment a continuous surveillance video temporally into activity patterns according to changes in video content.
3. Given a training video dataset, how to construct a model for the video content which can accommodate the variations in the unseen activity patterns both in terms of duration and temporal ordering?
4. Given a video content model and an unseen video, how to perform online activity recognition and unusual activity detection?

To this end, we first propose in this paper a semantic meaningful representation based on automatically detected discrete visual events in a video and two segmentation algorithms. One

* Corresponding author. Tel.: +44 20 7882 8020; fax: +44 20 8980 6533.

E-mail addresses: txiang@dcs.qmul.ac.uk (T. Xiang),
sgg@dcs.qmul.ac.uk (S. Gong).

of the proposed algorithms is offline offering better accuracy in segmentation but is computationally more demanding, while the other is online enabling real-time performance. We then develop a generative video content model based on unsupervised learning. Using this model an unseen activity pattern can be recognised into different classes if similar patterns are included in the training dataset and unusual activities can also be detected.

1.1. Video segmentation

A suitable representation is crucial for video segmentation. For activity based video segmentation, we propose to represent surveillance video content holistically in space and over time based on visual events detected automatically in the scene. This is in contrast to most previous approaches which represent video content at the signal level using image features such as colour, motion, and texture [1–6]. The essence of our method is to represent video content based on ‘what is happening’ rather than ‘what is present’ in the scene. ‘What is happening’ in the scene is reflected through activities captured in the video which most likely to involve multiple objects interacting or co-existing in a shared common space. An activity is composed of groups of co-occurring events which are defined as significant visual changes detected in image frames *over time*. Events are detected and classified by unsupervised clustering using Gaussian mixture model (GMM) with automatic model selection based on Schwarz’s Bayesian information criterion (BIC) [7,8]. Video content is then represented as temporally correlated events automatically labelled into different classes. By doing so changes in the presence of and temporal correlations among different classes of events can indicate video content changes therefore providing vital cues for activity based video segmentation.

Video segmentation has been studied extensively in the past two decades. Traditionally, a four-layer hierarchical structure is adopted for video structure analysis which consists of a frame layer, a shot layer, a scene layer and a video layer [4]. At the bottom of the structure, continuous image frames taken by a single camera are grouped into shots. A series of related shots are then grouped into a scene. Shot change detection is performed as the first step for video segmentation by most previous approaches [1–4]. The segmentation of shots and scenes heavily relies on a well-defined feature space usually dominated by colour and motion. For example, in Ref. [1], image frames were represented using a holistic colour histogram and the frame difference was exploited to detect shots. This structure is in general valid for constrained, well-structured broadcast videos of news and sports programmes. However, for a surveillance video which is taken continuously by a fixed camera *without* script-driven panning and zooming, global colour and motion information is either highly unreliable or unavailable [9]. More importantly, there is only one shot in a surveillance video and any shot-change detection based segmentation approach would be unsuitable.

Recently, DeMenthon et al. [10] proposed to represent a video as a high-dimensional temporal trajectory based on colour his-

togram and treat video segmentation as a trajectory breakpoint detection problem. Compared to the thresholding based segmentation algorithms adopted by most previous video segmentation approaches [1,2], a trajectory breakpoint detection based approach is more robust to local noise at individual frames because segmentation is performed holistically over the whole duration of the video. Various approaches have been proposed to segment a continuous trajectory into segments through breakpoint detection for time-series data segmentation [11–13]. Most of them are based on piecewise linear approximation (PLA) or probabilistic graphical models such as hidden Markov models (HMMs) [14]. PLA refers to finding the best approximation of a trajectory using straight lines by either linear interpolation or linear regression. However, the computational cost of PLA is nontrivial especially when the dimensionality of the trajectory space is high [13], resulting in most of the existing PLA segmentation algorithms only being applied to trajectories in a space with a dimensionality no bigger than 3. On the other hand, HMMs have the potential to be robust to noise and are capable of dynamic time warping (DTW). However, a large number of parameters are needed to describe an HMM when the dimensionality of the trajectory space is high. This makes an HMM vulnerable to over-fitting when training data are insufficient. To solve this problem, we propose a multi-observation hidden Markov model (MOHMM) which requires less parameters compared to a conventional HMM. It is thus more suitable for high-dimensional video content trajectory segmentation.

Most existing segmentation algorithms are offline. For the purpose of video segmentation, an online algorithm has its distinctive advantage due to the huge amount of surveillance video data to be processed and more importantly the real-time nature of some surveillance applications. For instance, if a video content change is detected online and in real-time, it can be used for alerting CCTV control room operators to act accordingly. One of the most popular online segmentation algorithms for temporal trajectories is the sliding window (SW) algorithm based on the SW principle [15] and local PLA. However, the SW algorithm tends to over-segment [11,15,16]. One possible explanation is that it lacks a global view of the data since it only ‘looks backward’ without ‘looking forward’. Keogh et al. [11] attempted to solve the problem by combining the bottom-up offline algorithm [15] with the SW principle. Nevertheless, their algorithm works only on trajectories with very short segments. It is thus impossible for the algorithm to run in real-time on a typical surveillance video sequence which comprises segments lasting over hours. In this paper, we propose a novel forward–backward relevance (FBR) algorithm. Compared to a conventional SW algorithm, FBR is less sensitive to noise and more importantly, can be run in real-time.

1.2. Video content modelling for activity recognition and unusual activity detection

Using the proposed segmentation algorithms, continuously recorded video or online CCTV input can be segmented into activity patterns. Given these activity patterns, the goal of video

content modelling is to learn a model that is capable of detecting unusual activity patterns whilst recognising novel instances of expected usual activity patterns. In this context, we define an unusual activity as an atypical activity pattern that is not represented by sufficient samples in a training dataset but critically it satisfies the specificity constraint to an unusual pattern. This is because one of the main challenges for the model is to differentiate unusual activity from outliers caused by noisy visual features used for activity representation. The effectiveness of an video content modelling approach shall be measured by (1) how well unusual activities can be detected (i.e. measuring specificity to expected patterns of activity) and (2) how accurately and robustly different classes of usual activity patterns can be recognised (i.e. maximising between-class discrimination).

To solve the problem, we develop a novel framework for fully unsupervised video content modelling and online unusual activity detection. Our framework has the following key components:

1. Discovering natural groupings of activity patterns using an activity affinity matrix. A number of affinity matrix based clustering techniques have been proposed recently [17–19]. However, these approaches require known number of clusters. Given an unlabelled dataset, the number of activity classes are unknown in our case. To automatically determine the number of clusters, a recently proposed spectral clustering algorithm [20] is deployed.
2. A composite generative video content model using a mixture of dynamic Bayesian networks (DBNs). The advantages of the such a generative video content model are twofold: (a) It can accommodate well the variations in the unseen and usual activity patterns both in terms of duration and temporal ordering by generalising from a training set of limited number of samples. This is important because in reality the same usual activity can be executed in many different usual ways. These variations cannot possibly be captured in a limited training dataset and need to be dealt with by a learned video content model. (b) Such a model is robust to errors in activity representation. This is because that a mixture of DBNs can cope with errors occurred at individual frames and is also able to distinguish an error corrupted usual activity pattern from an unusual one.
3. Online unusual activity detection using a run-time accumulative unusual activity measure and usual activity recognition using an online LRT method. A run-time accumulative measure is introduced to determine how usual/unusual an unseen activity pattern is on-the-fly. The activity pattern is then recognised as one of the usual activity classes if detected as being usual. Recognition of usual activities is carried out using an online LRT method which holds the decision on recognition until sufficient visual evidence has become available. This is in order to overcome any ambiguity among different activity classes observed online due to insufficient visual evidence at a given time instance. By doing so, robust activity recognition and unusual activity detection are ensured at the shortest possible time, as opposed to

previous work such as [9,21,22] which requires completed activity patterns being observed. Our online LRT based activity recognition approach is also advantageous over previous ones based on the *maximum likelihood* (ML) method [9,22,23]. An ML based approach makes a forced decision on activity recognition at each time instance without considering the reliability and sufficiency of the accumulated visual evidence. Consequently, it can be error prone.

The rest of the paper is structured as follows: in Section 2, we describe an event based surveillance video content representation approach. In Section 3, we address the problem of surveillance video segmentation. Two novel segmentation algorithms, MOHMM and FBR are introduced in Sections 3.1 and 3.2, respectively. Comparative experiments are conducted using over 10 h of challenging outdoor surveillance video footages and the results are presented in Section 3.3. The pros and cons of both algorithms are analysed. The advantage of our event based video content representation over the traditional image feature based representation is also made clear by our experimental results. The proposed video content modelling approach is described in Section 4, where experiments are also presented to evaluate the effectiveness and robustness of the approach. A conclusion is drawn in Section 5.

2. Semantic video content representation

We consider an activity based video content representation. Visual events are detected and classified automatically in the scene. The semantics of video content are considered to be best encoded in the occurrence of such events and the temporal correlations among them.

2.1. An example scenario

For clarity and concreteness, we shall illustrate our approach using surveillance videos monitoring aircraft docking operations at an airport. Around 15 different activities take place in the scene. Among them, only six activities are visually detectable. Fig. 1 shows an example activity based structure of a complete aircraft docking operation. One activity can be followed by either another activity immediately or a period of inactivity when no meaningful visual changes can be detected. The durations of activities and inactivity gaps vary hugely during a single aircraft docking operation and across different operations. There are also variations in the temporal order of activities and inactivity gaps. For example, the temporal order of *frontalCargoService* and *frontalCateringService* appears to be arbitrary. It is noted that some activities such as *airCraftArrival* involve the movement of a single object while other activities such as *frontalCargoService* and *frontalCateringService* consist of movements of multiple objects which may leave and re-appear in the scene. For the latter case, there often exist a number of short inactivity break-ups within an activity which are different from a long inactivity gap between

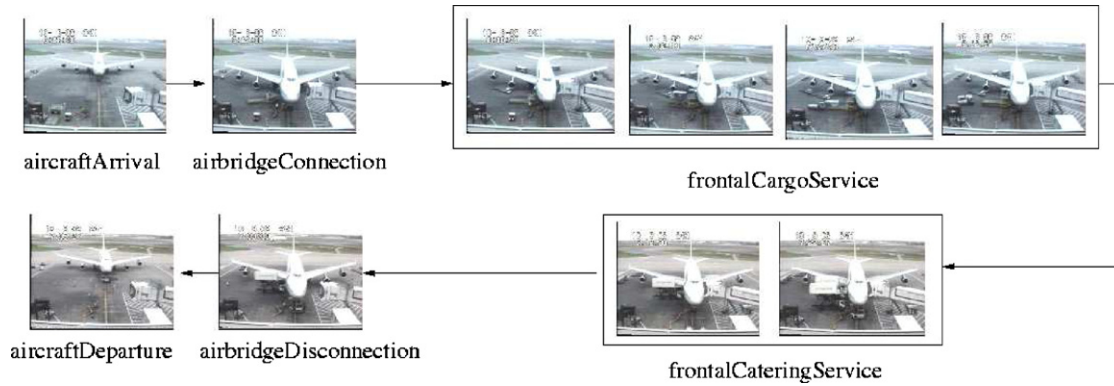


Fig. 1. An example structure of a complete aircraft docking video. Representative frames of different activities occurred in the scene are also shown.

two activities only by their durations. All these characteristics make analysis of surveillance video content very challenging.

2.2. Detecting and classifying visual events

We define events as significant scene changes characterised by the location, shape and direction of the changes. They are object-independent and location specific. We also consider that these events are autonomous, meaning that both the number of these events and their whereabouts in a scene are determined automatically bottom-up without top-down manual labelling using predefined hypotheses.

2.2.1. Seeding event: measuring pixel change history

Adaptive mixture background models are commonly used to memorise and maintain the background pixel distribution of a dynamic scene [24–26]. The major strength of such a model is its potential to cope with persistent movements of background objects such as waving tree leaves given appropriate model parameter setting. However, an adaptive mixture background model cannot differentiate, although may still be able to detect the presence of, pixel-level changes of different temporal scales. In general, a pixel-level change of different temporal scales can have different significance in its semantics:

1. A short term change is most likely to be caused by instant moving objects (e.g. passing-by people or vehicles).
2. A medium term change is most likely to be caused by the localised moving objects (e.g. a group of people standing and talking to each other).
3. A long term change is most likely to be caused by either the introduction of novel static objects into the scene, or the removal of existing objects from the scene (e.g. a piece of furniture is moved in the background or a car is parked in a carpark).

We seek a single, unified multi-scale temporal representation that can capture and differentiate changes of such different rates/scales at the pixel level. Temporal wavelets were adopted for such a multi-scale analysis [27]. However, the computational cost for multi-scale temporal wavelets at the pixel level

is very expensive. They are therefore unsuitable for real-time performance. Alternatively, motion history image (MHI) is less expensive to compute by keeping a history of temporal changes at each pixel location which then decays over time. MHI has been used to build holistic motion templates for the recognition of human movements [28] and moving object tracking [29]. An advantage of MHI is that although it is a representation of the history of pixel-level changes, only one previous frame needs to be stored. However, at each pixel location, explicit information about its past is also lost in MHI when current change are updated to the model with their corresponding MHI values ‘jumping’ to the maximal value. To overcome this problem, pixel signal energy was introduced to measure the mean magnitude of pixel-level temporal energy over a period of time defined by a backward window [25]. The size of the backward window determines the number of frames (history) to be stored. However, this approach suffers from its sensitivity to noise and also being expensive to compute.

Here we propose to use pixel change history (PCH) [30] for measuring multi-scale temporal changes at each pixel. The PCH of a pixel is defined as

$$P_{\zeta, \tau}(x, y, t) = \begin{cases} \min \left(P_{\zeta, \tau}(x, y, t-1) + \frac{255}{\zeta}, 255 \right) & \text{if } D(x, y, t) = 1, \\ \max \left(P_{\zeta, \tau}(x, y, t-1) - \frac{255}{\tau}, 0 \right) & \text{otherwise,} \end{cases} \quad (1)$$

where $P_{\zeta, \tau}(x, y, t)$ is the PCH for a pixel at (x, y) , $D(x, y, t)$ is a binary image indicating the foreground region, ζ is an accumulation factor and τ is a decay factor. When $D(x, y, t) = 1$, instead of jumping to the maximum value, the value of a PCH increases gradually according to the accumulation factor. When no significant pixel-level visual change is detected at a particular location (x, y) in the current frame, pixel (x, y) will be treated as part of the background and the corresponding PCH starts to decay. The speed of decay is controlled by a decay factor τ . The accumulation factor and the decay factor give us the flexibility of characterising pixel-level changes over time. In particular, large values of ζ and τ imply that the history of visual change at (x, y) is considered over a longer backward

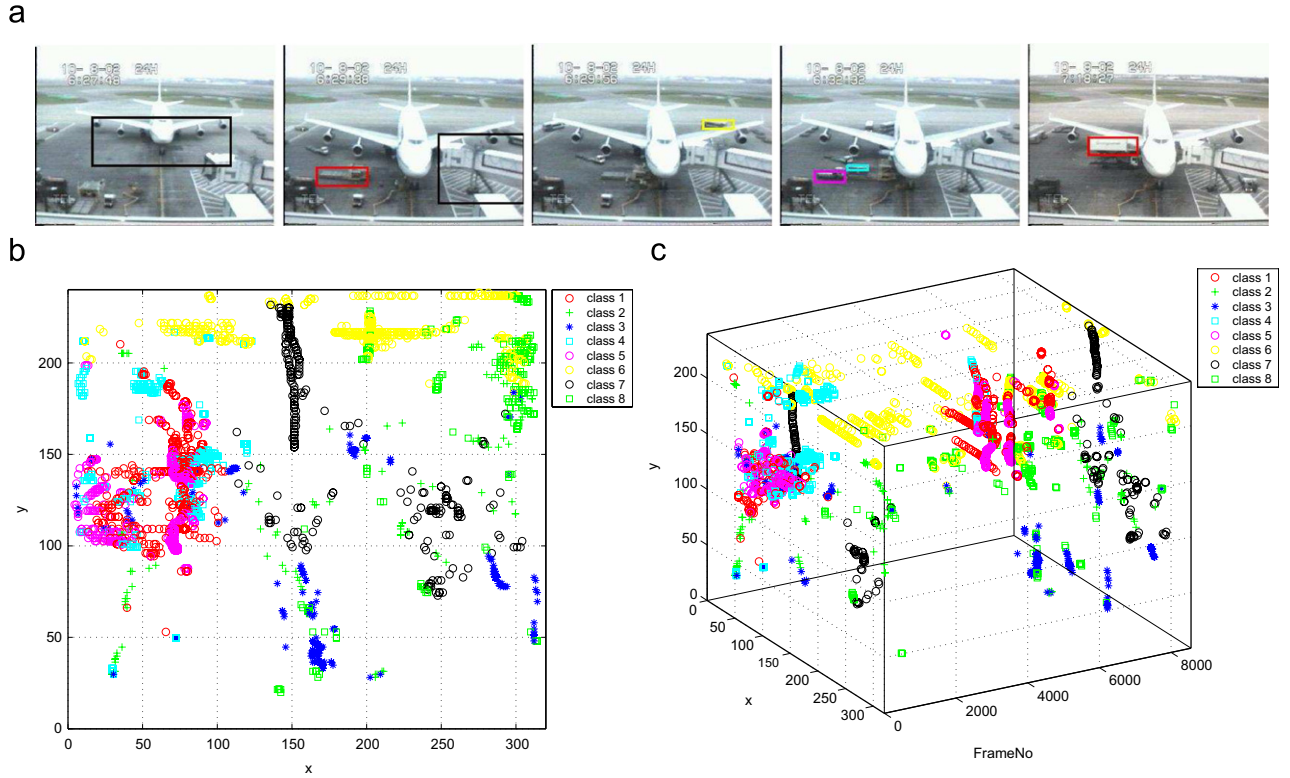


Fig. 2. Event detection and classification in an aircraft docking scene video. Eight classes of events are detected automatically, each of which corresponds to different movement patterns in the image frame. For example, event class 7 corresponds to the movements of the aircraft. Event classes 4–6 correspond to the movements of objects involved in activities *frontalCargoService* and *frontalCateringService*. (a) Events are detected and classified into different classes which are highlighted using bounding boxes in different colours. The spatial and temporal distribution of events of different classes are illustrated in (b) and (c), respectively, with centroids of different classes of events depicted using different colours.

temporal window. In the meantime, the ratio between ς and τ determines how much weight is put on the recent change.

If the binary image $D(x, y, t)$ in Eq. (1) is determined by the temporal difference between the current frame and the dynamic background maintained by an adaptive mixture model, a PCH based foreground model can be introduced to detect the medium and long term pixel changes. Specifically, we detect those pixels that are associated with medium term changes by the following condition:

$$|I(x, y, t) - I(x, y, t - 1)| > T_M, \quad (2)$$

where T_M is a threshold. Pixel level changes that do not satisfy the above condition are caused by long term changes such as the introduction of static novel objects into the scene or the removal of existing objects from the scene. Note that Eq. (2) is used for distinguishing the detected foreground pixels according to the nature of the visual changes. It is not used for detecting foreground pixels using background subtraction.

2.2.2. From pixel groups to unsupervised clustering and classification of events

Given detected pixel changes in each image frame, we aim to form discrete events. The connected component method is adopted to group those changed pixels. Small groups are then removed by a size filter and the rest groups with an average

PCH (of the PCHs for all the pixels within each group) larger than a threshold T_B are referred to as salient pixel groups and considered as events. An event is represented by a 7-D feature vector

$$\mathbf{v} = [\bar{x}, \bar{y}, w, h, R_m, M_px, M_py], \quad (3)$$

where (\bar{x}, \bar{y}) is the centroid of the salient pixel group, (w, h) are the width and height of the salient pixel group, R_m represents the percentage of those pixels in the group that satisfy Condition (2), and (M_px, M_py) are a pair of first-order moments of the PCH image within the salient pixel group. Among these features, (\bar{x}, \bar{y}) are location features, (w, h) are shape features, R_m is visual change type feature and (M_px, M_py) are motion features capturing the direction of object motion direction.¹ Note that in our approach, salient pixel groups are defined within each image frame. Alternatively, salient groups can be defined in a spatio-temporal volume which could in theory lead to better clustering. One could adopt a method such as the one proposed by Greenspan et al. [31]. Alternatively, we have also developed an approach for salient event detection over a spatio-temporal volume using multi-scale entropy ratio

¹ Similar to the MHI (see Ref. [28]), PCH implicitly represents the direction of movement. First-order moments based on PCH value distribution within the bounding box is thus capable of measuring the direction of movement quantitatively.




Fig. 3. The example video shown in Fig. 2 is represented by scene vectors evolving over time. We have $K_e = 8$ for this video. s_t^k is depicted by a dot in colour when $s_t^k > 0$. There are frequent but short inactivity break-ups within activities (see between frames 1 and 2000) and a long inactivity gaps between activities (between frames 3000 and 6000).




Fig. 4. The eight elements of $\tilde{\mathbf{s}}_t$ over time for the example video shown in Fig. 2. Each element of $\tilde{\mathbf{s}}_t$ is normalised to have a value range of $[0, 1]$.

over space and time presented elsewhere [32]. However, such a spatio-temporal volume based events detection and recognition approach is always computationally expansive, and may not be tractable given the complexity of the activities captured in video footages. In order to achieve real-time performance, we decided to make a compromise between the speed and performance of the event detection and recognition algorithm by defining the salient groups within each image frame.

Salient pixel groups are clustered and classified unsupervised into different events in the 7-D feature space using a GMM. The GMM is estimated using expectation-maximisation (EM) [33] and the model order of the GMM is determined using the BIC [7]. An example of event detection and classification is shown in Fig. 2.

2.3. Constructing a scene vector

Classified events can be considered as ‘snapshots’ of activities captured in a scene. To hide the image feature information

and focus on the semantic video content, a scene vector is constructed for each image frame of a video. A scene vector \mathbf{sv}_t for a video frame t is defined as

$$\mathbf{sv}_t = [s_t^1, \dots, s_t^k, \dots, s_t^{K_e}], \quad (4)$$

where K_e is the number of event classes automatically determined using BIC. The value of s_t^k is the number of events of the k th event class detected in the frame t .

A scene vector gives a description of ‘what is happening’ in the scene through the class labels of the detected events. It is thus a concise representation of the video content at the semantic level (see Fig. 3 for an example). However, directly using this scene vector to detect video content changes can cause problems for video segmentation. Specifically, the value of a scene vector \mathbf{sv}_t can become $\mathbf{0}$ (i.e. absent of any event at a given frame) frequently throughout the video sequence (see Fig. 3). This can be caused by either frequent but short inactivity break-ups within activities or long inactivities between activities. Each ‘coming to zero’ is reflected as a dramatic change in

the value of \mathbf{sv}_t due to the discrete nature of s_t^k . Those changes that correspond to real changes of video content can thus easily be overwhelmed by changes caused by the inactivity break-ups within activities, which makes temporal segmentation of the video difficult.

2.4. Representing video content over time

To overcome this problem, let us now consider representing the video content using a cumulative scene vector computed at frame t using \mathbf{sv}_t from frame 1 to the frame t . More specifically, the k th element of the cumulative scene vector (denoted as $\tilde{\mathbf{sv}}_t$) is computed as

$$\tilde{s}_t^k = \sum_{i=1}^t s_i^k. \quad (5)$$

The value of each element of $\tilde{\mathbf{sv}}_t$ will increase monotonically with time (see Fig. 4). Compared to the scene vector representation \mathbf{sv}_t (see Fig. 3), the short inactivity break-ups at individual frames have little impact on the values of the cumulative scene vector evolved over time. It thus becomes easier to detect breakpoints that correspond to significant changes in video content.

3. Temporal segmentation of surveillance videos

It has been shown in the preceding section that a cumulative scene vector can represent the video content at the semantic level over time and is capable of capturing video content changes despite variations in activity durations and occurrences of inactivity break-ups within activities. After mapping a video sequence into a cumulative scene vector trajectory,² the breakpoints on the trajectory correspond to the video content change points. We thus consider the video segmentation problem as a temporal video content trajectory breakpoint detection problem. To solve the problem, one offline and one online algorithms are proposed in this section.

3.1. Multi-observation hidden Markov model (MOHMM)

We first consider detecting breakpoints on a temporal video content trajectory using a DBN, more specifically, an MOHMM. The DBNs are Bayesian belief networks (BBNs³) that have been extended to model time series data [34,35]. A DBN \mathbf{B} is described by two sets of parameters (\mathbf{m}, Θ). The first set \mathbf{m} represent the structure of the DBN which include the number of hidden variables and observation variables per time instance, the number of states for each hidden variable and the topology of the network (set of directed arcs connecting nodes). The i th hidden variable and the j th observation variable

at time instance t are denoted as $S_t^{(i)}$ and $O_t^{(j)}$, respectively, where $i \in \{1, \dots, N_h\}$ and $j \in \{1, \dots, N_o\}$ and N_h and N_o are the number of hidden variables and observation variables, respectively. The second set of parameters Θ quantify the state transition models $P(S_t^{(i)}|Pa(S_t^{(i)}))$, the observation models $P(O_t^{(j)}|Pa(O_t^{(j)}))$ and the initial state distributions $P(S_1^{(i)})$ where $Pa(S_t^{(i)})$ are the parents of $S_t^{(i)}$ and similarly, $Pa(O_t^{(j)})$ for observations. In this paper, unless otherwise stated, $S_t^{(i)}$ are discrete and $O_t^{(j)}$ are continuous random variables. Each observation variable has only hidden variables as parents and the conditional probability distributions (CPDs) of each observation variable are Gaussian for each state of its parent nodes.

When a DBN is adopted for video content modelling, the observation variables in the network correspond to the cumulative scene vector $\tilde{\mathbf{sv}}_t$ and hidden states in the network correspond to the video content. After model parameter estimation and hidden state inference, changes in hidden states should reflect the video content changes and thus can be used for video temporal segmentation. DBNs of different topologies can be employed for video structure modelling. We can choose the DBNs with the simplest topology, the HMMs (Fig. 5(a)). In an HMM, the observation variable at each time instance corresponds to $\tilde{\mathbf{sv}}_t$, which is of dimension K , i.e. the number of event classes. However, a drawback of using HMMs for modelling complex temporal process is that a large number of parameters are needed to describe the model when the observation variables are of high dimension. This may result in poor model learning given insufficient training data. To address this problem, various topological extensions to the standard HMMs can be considered to factorise the state and/or observation space by introducing multiple hidden variables and multiple observation variables. In our case, we need to detect video content changes based on hidden state changes. Because the hidden states of the single hidden variable at each time instance corresponds clearly to the video content segments, the factorisation in the states space is unnecessary. Therefore, we only perform factorisation in the observation space which results in a MOHMM shown in Fig. 5(b). Specifically, $O_t^{(j)}$ ($j \in \{1, \dots, K\}$) in Fig. 5(b) is a 1-D random variable corresponding to one of the K elements of $\tilde{\mathbf{sv}}_t$. This is to factorise the observational space based on the assumption that the occurrence of events of different event classes is independent from each other. Consequently, the number of parameters for describing a MOHMM is much lower than that for an HMM ($2KN_h + N_h^2 - 1$ for an MOHMM and $(K^2 + 3K)N_h/2 + N_h^2 - 1$ for an HMM).

The remaining problem is to automatically determine N_h , the number of hidden states which corresponds to the number of video segments. To this end, we adopt Schwarz's BIC [7] to automatically determine N_h . For a model \mathbf{m}_i parameterised by a C_i -dimensional vector $\Theta_{\mathbf{m}_i}$, the BIC is defined as

$$BIC = -2 \log L(\Theta_{\mathbf{m}_i}) + C_i \log N, \quad (6)$$

where $L(\Theta_{\mathbf{m}_i})$ is the maximal likelihood under \mathbf{m}_i , C_i is the dimensionality of the parameters of \mathbf{m}_i and N is the size of the

² More precisely, it is a video polyline due to the discrete nature of the image frames.

³ BBNs are also known as Bayesian networks, belief networks or directed acyclic graphical (DAG) models. They are special cases of graphical models which combine probability theory and graph theory to address two important issues in data modelling: uncertainty and complexity.




Fig. 5. (a) HMM and (b) MOHMM with the observation nodes illustrated as shaded circles and hidden nodes as clear circles.

dataset. For the MOHMM, $L(\Theta_{\mathbf{m}_i})$ can be written as

$$-2 \log \left\{ \sum_{S_t} \left\{ P(S_1) \prod_{t=2}^T P(S_t | Pa(S_t)) \right. \right. \\ \left. \left. \times \prod_{t=1}^T \prod_{j=1}^K P(O_t^{(j)} | Pa(O_t^{(j)})) \right\} \right\},$$

where $Pa(S_t)$ are the parents of S_t and similarly, $Pa(O_t^{(i)})$ for observations. The search for the optimal N_h that produces the minimal BIC value involves parameter learning. More specifically, for each candidate state number, the corresponding parameters are learned iteratively using the EM algorithm. The E step, which involves the inference of hidden states given the parameters estimated in the last M step, can be implemented using an exact inference algorithm such as the junction tree algorithm [36]. After parameter learning, the BIC value can be computed using Eq. (6) where $L(\Theta_{\mathbf{m}_i})$ has been obtained from the final M step of EM for parameter learning.

3.2. Forward–backward relevance (FBR)

The inference and learning of a MOHMM is computational expensive especially when a video content trajectory is long and of a high dimension. For computational efficiency and real-time performance, a novel online video segmentation algorithm is proposed which is based on an FBR principle. More specifically, at each time instance, the relevance of one vertex on the video trajectory is measured with respect to the starting point of the current video segment (looking backward) and a certain distance ahead on the trajectory (looking forward). The relevance of this vertex is minimal when the three vertices sit on a straight line in the high-dimensional trajectory space. When the video trajectory turns drastically at the vertex, indicating a video content change, a high relevance value will be obtained. This relevance measure can thus be used for breakpoint detection on a video trajectory. Our algorithm is motivated by the offline discrete curve evolution (DCE) segmentation algorithm and aims to overcome the over-segment tendency of the conventional online SW algorithm.

DCE was originally proposed for 2-D shape decomposition [12]. It can readily be extended for detecting breakpoints of high-dimensional video content trajectories. The algorithm starts with the finest possible segmentation (i.e. $T - 1$ segments

```

input : A  $K$  dimensional trajectory of length  $T$ :  $\bar{s}\mathbf{v}_t$ 
output: Detected breakpoints on the trajectory:  $breakPoint[]$ 

1  $breakpointNo = 0$ ;
2 while  $t < T$  do
3   if  $t > (breakPoint[breakpointNo] + L_{min})$  then
4     compute the relevance of the  $t - L_{min}/2$ -th vertex as:
        $R(t - L_{min}/2) = R(breakPoint[breakpointNo], t - L_{min}/2, t)$ ;
5     if  $R(t - L_{min}/2) > Th$  then
6        $breakpointNo = breakpointNo + 1$ ;
7        $breakPoint[breakpointNo] = t - L_{min}/2$ ;
8     end
9   end
10 end

```

Fig. 6. An online video segmentation algorithm based on FBR. In line 4, $R(t - L_{min}/2)$ is computed using Eq. (7) with the function $d(\cdot)$ defined as the Euclidean distance.

for a trajectory of length T) and progresses with two segments being merged at each step. Specifically, the cost of merging all adjacent pairs of segments are computed and the pair with the minimal cost is merged. This process is performed iteratively until a stopping criterion is met. DCE computes the merging cost as a relevance measure R which is computed for every vertex v and depends on v and its two neighbour vertices u and w :

$$R(u, v, w) = d(v, u) + d(v, w) - d(u, w), \quad (7)$$

where function $d(\cdot)$ can be any distance or similarity measures such as the Euclidean distance. In each iteration, the vertex with minimal R is deleted until the minimal R exceeds certain threshold or only a user-specified number of vertices are left in the trajectory. The remaining vertices correspond to the breakpoints on the video trajectory. Although DCE is similar in spirit to the bottom-up segmentation algorithm based on PLA in time series data analysis [11], it is much faster compared to the bottom-up algorithm using either linear interpolation or linear regression, especially when the dimensionality of the trajectory space is high.

In our online segmentation algorithm, instead of computing the relevance of each vertex iteratively, the relevance of only one vertex is computed at each time instance when a new image frame has arrived. The algorithm is outlined in Fig. 6.

Two parameters, L_{min} and Th need to be determined. In this paper, they are learned from a training dataset. Specifically, we perform DCE on the training data and L_{min} is set to half of the length of the shortest segment in the training data and Th is computed as the average relevance value of the breakpoints detected by DCE.

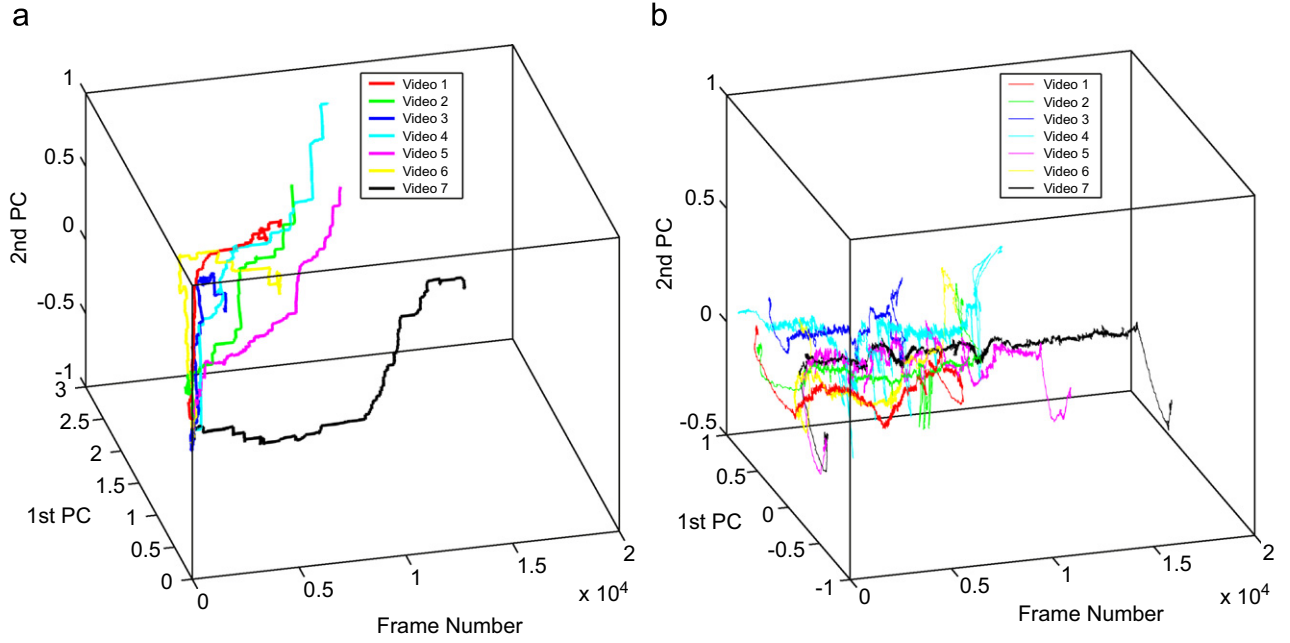


Fig. 7. (a) and (b) show the first two principal components of $\tilde{\mathbf{sv}}_t$ and colour histogram evolving over time for the seven aircraft docking videos, respectively. Note that PCA was used only for visualisation here. Each element of $\tilde{\mathbf{sv}}_t$ and colour histogram is normalised to have a value range of $[0, 1]$.

There are a number of important characteristics of the algorithm that distinguish it from previous ones:

1. For each new data point, instead of fitting the trajectory to a straight line, FBR computes a relevance measure of a single vertex which involves computing only three Euclidean distances. The computation complexity is *independent* of the segment length and much lower than the SW algorithm and its variations.
2. At each time instance, only the coordinates of $L_{\min}/2 + 1$ vertices need to be stored for computation. This makes our algorithm suitable for long videos such as 24/7 surveillance videos where it could be hours before any video content changes take place.
3. Compared to the SW algorithm and its variations, our FBR algorithm is more robust to noise because it has a global view of the trajectory by looking both forward and backward.
4. The FBR algorithm has a $L_{\min}/2$ delay in detecting break-points which in practice is a very small number compared to the length of a typical video segment. This is the price one has to pay to achieve robustness in detection.

3.3. Experiments

Experiments were conducted on the representation and segmentation of CCTV surveillance videos monitoring an aircraft ramp area. A fixed CCTV analogue camera took continuous recordings. After digitisation from VHS tapes, the final video sequences have a frame rate of 2 Hz. Note that it is not uncommon to have such an extremely low frame rate for CCTV surveillance videos, which makes the video segmentation prob-

lem even more challenging. Each image frame has a size of 320×240 pixels. Our database for the experiments consists of seven videos of aircraft docking lasting from 6470 to 17 262 frames per sequence (around 50–140 min of recording), giving in total 72 776 frames of video data (around 10 h in duration). These video data cover different times of different days under changing lighting conditions, from early morning, midday to late afternoon. They are referred to as video 1 to video 7, respectively. Among the seven videos, videos 1–6 follow the typical video structure with cargo services performed before catering services while video 7 has a different video structure with cargo services performed after catering services. In the following we present results on (1) comparative performance evaluation on video segmentation using the offline MOHMM and DCE algorithms, and the online FBR algorithm; and (2) comparing our cumulative scene vector based video content representation with a colour histogram based one.

3.3.1. Video content representation

Eight classes of events were automatically detected and classified from the aircraft docking video sequences. They were caused by the movements of the docking aircraft and various ground vehicles such as catering trucks, cargo trucks and cargo lifts (see Fig. 2 for an example). A cumulative scene vector $\tilde{\mathbf{sv}}_t$ (Eq. (5)) was used to represent the video content. Fig. 7(a) shows the evolution of the first two principal components of $\tilde{\mathbf{sv}}_t$ for the seven different sequences.⁴ It can be seen that the video content trajectory of video 7 is distinctively different from

⁴ Note that principal component analysis (PCA) was performed only for visualisation of the high-dimensional video trajectory. Video segmentation was carried out in the original 8-D video trajectory space.

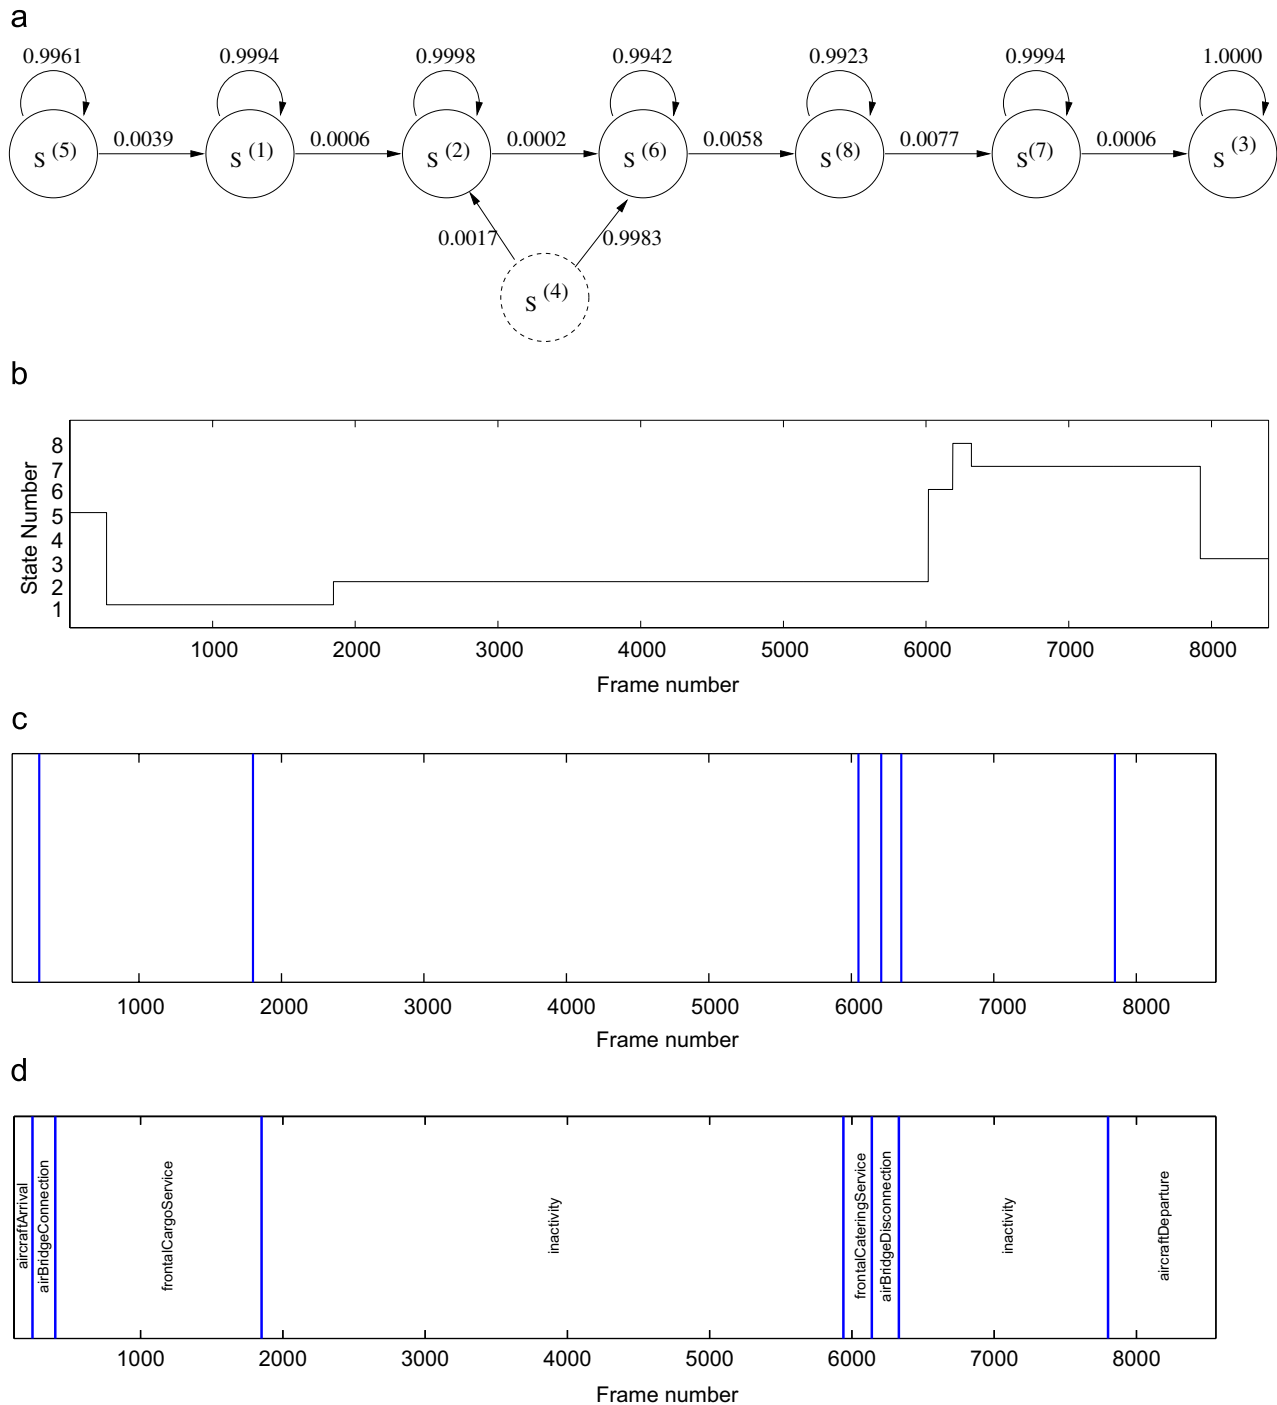


Fig. 8. Temporal segmentation of video 1 using MOHMM. In (a), the inferred hidden state 4 is a virtual state that does not model data (the probability of staying in state 4 is zero). (a) Transition probabilities of hidden states of the MOHMM trained for video 1. (b) The inferred states of MOHMM. (c) MOHMM segmentation. (d) Manual segmentation.

those of the other six videos. For comparison, the same video sequences were also represented using colour histograms. We defined four histogram bins for each of the three colour channels of the RGB colour space. Separate RGB colour histograms were concatenated into a single one with 12 components in each frame. Fig. 7(b) illustrates the evolution of the first two principal components of the colour histogram representation over time from the same seven sequences as Fig. 7(a). It is evident

from Fig. 7(b) that this colour histogram based video content trajectories are less smooth compared to the video content trajectories based on the cumulative scene vector. It is also noted that only the arrival and departure of the aircraft are captured well by the colour histogram changes (see the beginning and ending part of the trajectories). It can be seen from Fig. 7(b) that the trajectory corresponding to video 7 is similar to other trajectories under the colour histogram representation.

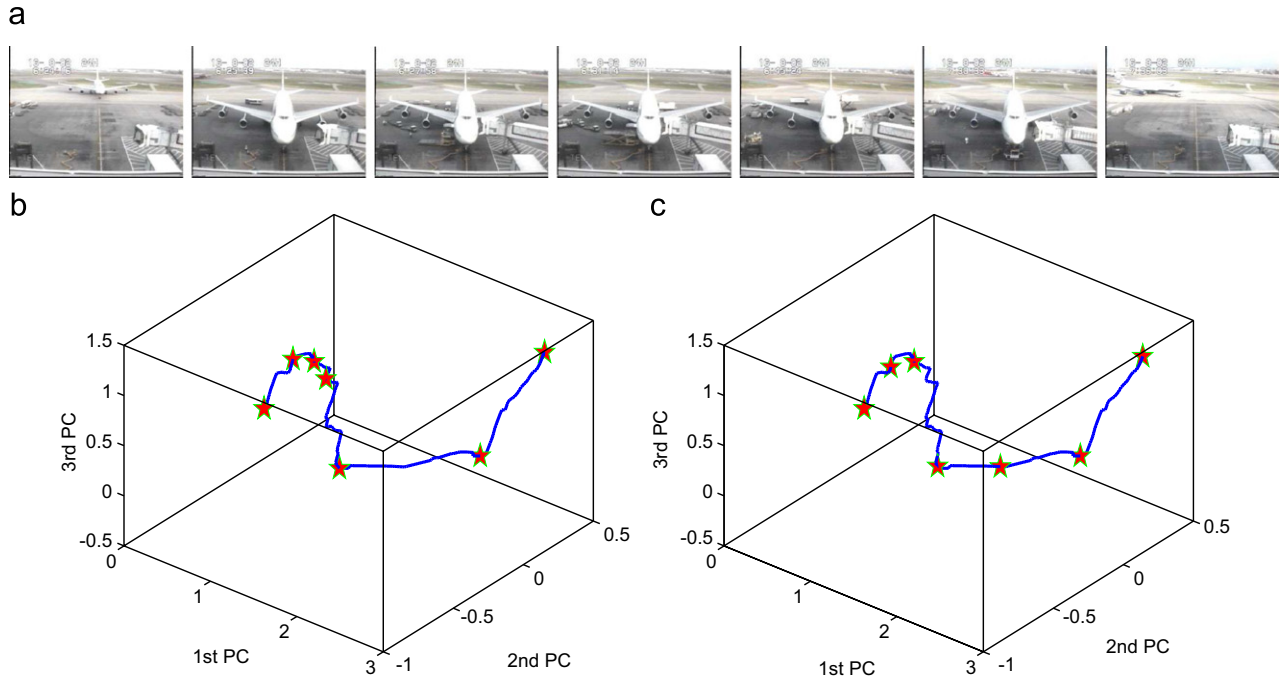


Fig. 9. An example breakpoint detection result by the FBR algorithm. (a) The frames corresponding to the detected breakpoints. (b) Detected breakpoints overlapped on the video content trajectory. (c) Manual segmentation result. Notice that although the video are illustrated in a 3-D PCA space in (b) and (c), it was segmented in an 8-D video trajectory space.

3.3.2. Video segmentation

The seven videos were first manually segmented into activities and inactivity gaps to give the ground truth of the breakpoints. The number of segments in each sequence was between 6 and 11. The videos were then divided into training and test sets with three videos in the training set and four in the testing set. This was repeated 10 times and in each trial a different training set of three videos were randomly chosen in order to avoid any bias in comparing the performance of different segmentation algorithms. We compared the performance of the offline DCE and MOHMM algorithms and our online FBR algorithm on the videos represented by scene vector $\tilde{\mathbf{s}}_t$ (Fig. 7(a)). Note that in our experiments, training data were needed by both DCE and FBR for determining the number of segments. However, they are not required directly by MOHMM because it automatically determines the number of video segments through model selection (see Section 3.1).

An example of video segmentation using MOHMM is shown in Fig. 8. A MOHMM was trained for video 1. The number of observation variables in each time slice was 8, which was automatically determined by the number of event classes. The number of hidden states was automatically determined to be 8 using the BIC. All the model parameters are initialised randomly. The learned state transition probabilities are illustrated in Fig. 8(a). The inferred states of MOHMM are shown in Fig. 8(b). The video was then segmented based on changes on the inferred states (see Fig. 8(c)). A manual segmentation with a short description of activities occurred in each segment is shown in Fig. 8(d). Comparing Fig. 8(c) with Fig. 8(d), we can see that most video content changes have been detected correctly. It

is noted that a breakpoint between aircraftArrival and airbridgeConnection was not detected. This was due to the overlapping in time between this two activities. It is also observed that some sparsely and randomly appearing events corresponding to a number of passing-by vehicles in the aircraft ramp area (see Fig. 3 between frames 2000 and 6000) were not regarded as being significant and were thus not picked up by our model as video content changes, thanks to the cumulative scene vector based video content representation.

Fig. 9 shows an example of breakpoint detection using the FBR algorithm. Fig. 9(b) shows video 2 represented by the cumulative scene vector $\tilde{\mathbf{s}}_t$ (illustrated in a 3-D PCA space) with the seven detected breakpoints overlapped on the video content trajectory (the trajectory evolves from left to right). The top row of Fig. 9 shows the frames corresponding to the detected breakpoints. Compared with the manual segmentation result (Fig. 9(c)), it can be seen that the fourth breakpoint is a false positive which was caused by a long inactivity break-up during the frontal cargo service. A breakpoint between the fifth and sixth detected breakpoints was missed which corresponds the ending point of a frontal catering service. This was caused by the quick departure of the catering truck and the low frame rate of the video.

The performance of the three algorithms over 10 trials was measured by the detection rate and false positive rate of breakpoint detection. A breakpoint detected within 100 frames of the true breakpoint was regarded as being correctly detected. In total, there were 334 true breakpoints in the testing sets over 10 trials. The distance measure in Eq. (7) was Euclidean for both DCE and FBR. Table 1 shows that satisfactory results

Table 1
Segmentation results on cumulative scene vector trajectories

	DCE	MOHMM	FBR
Detection rate (%)	78.1	82.3	75.6
False positives (%)	0.0236	0.0188	0.0388

Table 2
Comparing the computational costs of the three segmentation algorithms

	DCE	MOHMM	FBR
Computational cost (second per frame)	0.023	0.098	0.000032

Table 3
Segmentation results on colour histogram trajectories

	DCE	MOHMM	FBR
Detection rate (%)	31.4	26.7	28.4
False positives (%)	0.0624	0.0430	0.0539

were obtained by all three algorithms considering the challenging nature of the problem and the noisy data used in the experiments. The FBR algorithm achieved slightly worse detection rate compared to the two offline algorithms and tended to over-segment the video sequences. It is also noted that the starting and ending points of the frontal catering service were often miss-detected by all three algorithms. In particular, when a catering truck moved into position very quickly during the frontal catering service, the resultant change in the video content trajectory was similar to those caused by passing vehicles which were not involved in any activity. This type of changes would not be picked up by the segmentation algorithms. On the other hand, the long inactivity break-ups caused by the lack of movements during the frontal catering service were the main reason for false positives. These issues could be partially addressed by employing a more sophisticated background model.

The algorithms were implemented on a Xeon 3 GHz platform. The computational cost for background modelling and event detection was 0.11 s per frame. Given the constructed cumulative scene vector, the computational cost of the three segmentation algorithms are compared in Table 2. It can be seen that the computational cost for FBR is much low than those of the other two. Since FBR is an online algorithm, our online segmentation algorithm can run at a near real-time speed (9 Hz) without any code optimisation (i.e. the speed is mainly determined by the background modelling and event detection).

3.3.3. Comparison with colour histogram based representation

We also performed DCE, MOHMM and FBR on the colour histogram based video content representation following the same experimental procedure. It is shown in Table 3 that all three algorithms yielded much inferior results on the same seven aircraft docking videos. In particular, it is noted that only the arrival and departure of the aircraft were detected reliably. The changes caused by the start and end of other activities were

easily overwhelmed by image noise and illumination changes frequently encountered in a typical outdoor environment. This resulted in the lower detection rates and higher false positive rates (comparing Table 3 with Table 1).

3.3.4. Discussions

The key findings of our experiments are summarised and discussed as below:

1. Comparing the two proposed segmentation algorithms, MOHMM achieved higher detection rate and lower false alarm rate. Using MOHMM, video content is modelled explicitly by hidden states which need to be inferred from the observations. This makes MOHMM more robust to representation noise. However, it is also computationally more expensive than FBR. Also importantly, FBR is an online algorithm that can run in real-time together the background modelling and event detection. This makes FBR attractive for applications for which real-time performance is the key priority.
2. Our results show that a semantically meaningful video representation is crucial for segmentation of surveillance videos. This is related to the problem of ‘bridging the semantic gap’ which has received much attention recently in multimedia content analysis [5,37,38]. In our case, there is a gap between the low level image features (e.g. colour and motion) and semantic meaningful segmentation results (i.e. each segment corresponds to an activity which can be described using semantic concepts). To bridge this gap, we introduce the event-based video content representation with each class of events being constituents of different activities. When a new activity takes place, the changes in the event classes or/and the temporal correlations among them can be picked up by our segmentation algorithm for video content breakpoint detection.
3. Constructing a scene vector \mathbf{sv}_t is essentially a vector quantisation process which is useful for concise representation. However, using the scene vector directly is not appropriate for segmentation because of the discreteness of the scene vector. The introduction of the cumulative scene vector $\tilde{\mathbf{sv}}_t$ is critical to make the representation less sensitive to short inactivity break-ups within activities and noise in event detection.

4. Video content modelling for activity recognition and unusual activity detection

Using the approaches described in the preceding section, continuous surveillance videos in a training dataset are segmented into N video segments, ideally each video segment now containing a single activity. Taking an unsupervised approach, the task now is to discover the natural grouping of the activity patterns (i.e. clustering), upon which a video content model is built for the N video segments. Given an unseen activity, the model is then employed to detect if the activity pattern is different from what has been seen before (i.e. an unusual activity), and

recognise it into different activity classes if a similar pattern has been observed previously.

4.1. Activity representation

The activity pattern captured in the n th video segment \mathbf{v}_n is represented as a feature vector \mathbf{P}_n , given as

$$\mathbf{P}_n = [\mathbf{p}_{n1}, \dots, \mathbf{p}_{nt}, \dots, \mathbf{p}_{nT_n}], \quad (8)$$

where T_n is the length of the n th video segment and the t th element of \mathbf{P}_n is a K_e dimensional variable:

$$\mathbf{p}_{nt} = [p_{nt}^1, \dots, p_{nt}^k, \dots, p_{nt}^{K_e}]. \quad (9)$$

\mathbf{p}_{nt} corresponds to the t th image frame of \mathbf{v}_n where p_{nt}^k is the posterior probability that an event of the k th event class has occurred in the frame given the learned GMM. If an event of the k th class is detected in the t th image frame of \mathbf{v}_n , we have $0 < p_{nt}^k \leq 1$; otherwise, we have $p_{nt}^k = 0$.

It is worth pointing out that: (1) An activity pattern is decomposed into temporally ordered, semantically meaningful scene-events. Consequently, the activity representation is compact and concise. This is critical for a model-based video content modelling approach because model construction based upon concise representation is more likely to be computationally tractable for complex activities. (2) Different types of activity patterns can differ either in the classes of events they are composed of, or in the temporal orders of the event occurrence. For instance, activity patterns A and B are deemed as being different if (1) A is composed of events of classes a, b and d, while B is composed of events of classes a, c and e; or (2) Both A and B are composed of events of classes a, c and d; however, in A, event (class) a is followed by c, while in B, event (class) a is followed by d.

4.2. Activity clustering

The video content modelling problem can now be defined formally. Consider a training dataset \mathbf{D} consisting of N feature vectors:

$$\mathbf{D} = \{\mathbf{P}_1, \dots, \mathbf{P}_n, \dots, \mathbf{P}_N\}, \quad (10)$$

where \mathbf{P}_n is defined in Eq. (8) representing the activity pattern captured by the n th video segment. The problem to be addressed is to discover the natural grouping of the training activity patterns upon which a model for usual activity can be built. This is essentially a data clustering problem with the number of clusters unknown. There are a number of aspects that make this problem challenging: (1) Each feature vector \mathbf{P}_n can be of different lengths. Conventional clustering approaches such as K -means and mixture models require that each data sample is represented as a fixed length feature vector. These approaches thus cannot be applied directly. (2) A definition of a distance/affinity metric among these variable length feature vectors is nontrivial. Measuring affinity between feature vectors

of variable length often involves DTW [39]. A standard DTW method used in computer vision community would attempt to treat the feature vector \mathbf{P}_n as a K_e dimensional trajectory and measure the distance of two activity patterns by finding correspondence between discrete vertices on two trajectories. Since in our framework, an activity pattern is represented as a set of temporal correlated events, i.e. a stochastic process, a stochastic modelling based approach is more appropriate for distance measuring. (3) Model selection needs to be performed to determine the number of clusters. To overcome the above-mentioned difficulties, a recently proposed spectral clustering algorithm with feature and model selection is deployed [20].

To apply a spectral clustering algorithm, a data affinity matrix needs to be constructed. Here we only describe the procedure for constructing an affinity matrix for activity clustering and leave out the details of the spectral clustering algorithm which can be found in Ref. [20]. We utilise DBNs to provide a dynamic representation of each activity pattern feature vector in order to both address the need for dynamic warping and provide a string similarity metric. More specifically, each activity pattern in the training set is modelled using a DBN. To measure the affinity between two activity patterns represented as \mathbf{P}_i and \mathbf{P}_j , two DBNs denoted as \mathbf{B}_i and \mathbf{B}_j are trained on \mathbf{P}_i and \mathbf{P}_j , respectively, using the EM algorithm [34,40]. The affinity between \mathbf{P}_i and \mathbf{P}_j is then computed as

$$A_{ij} = \frac{1}{2} \left\{ \frac{1}{T_j} \log P(\mathbf{P}_j | \mathbf{B}_i) + \frac{1}{T_i} \log P(\mathbf{P}_i | \mathbf{B}_j) \right\}, \quad (11)$$

where $P(\mathbf{P}_j | \mathbf{B}_i)$ is the likelihood of observing \mathbf{P}_j given \mathbf{B}_i , and T_i and T_j are the lengths of \mathbf{P}_i and \mathbf{P}_j , respectively.

DBNs of different topologies can be used. However, it is worth pointing out that since a DBN needs to be learned for every single activity pattern in the training dataset which could be short in duration, a DBN with less number of parameters is more desirable. To this end, we employ the MOHMM described in Section 3.1. Note that when applying MOHMM for modelling each activity pattern, the number of hidden states for the MOHMM is set to K_e , i.e. the number of event classes. This is reasonable because the value of N_s should reflect the complexity of an activity pattern, so should the value of K_e .

4.3. Modelling video content using a composite model

Using the spectral clustering algorithm proposed in Ref. [20], the N activity patterns in the training dataset is grouped into K clusters with K being automatically determined via model order selection. To build a model for the observed/expected activities, we first model the k th activity class using a MOHMM \mathbf{B}_k . The parameters of \mathbf{B}_k , $\theta_{\mathbf{B}_k}$ are estimated using all the patterns in the training set that belong to the k th class. An activity model \mathbf{M} is then formulated as a mixture of the K MOHMMs. Given an unseen activity pattern, represented as an activity pattern feature vector \mathbf{P} as described in Section 4.1, the likelihood of

observing \mathbf{P} given \mathbf{M} is

$$P(\mathbf{P}|\mathbf{M}) = \sum_{k=1}^K \frac{N_k}{N} P(\mathbf{P}|\mathbf{B}_k), \quad (12)$$

where N is the total number of training activity patterns and N_k is the number of patterns that belong to the k th activity class.

4.4. Online activity recognition and unusual activity detection

Once constructed, the composite video content model \mathbf{M} can be used to detect whether an unseen activity pattern is usual using a run-time unusual activity measure. If detected as being usual, the activity pattern is then recognised as one of the K classes of usual activity patterns using an online likelihood ratio test (LRT) method.

An unseen activity pattern of length T is represented as $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_t, \dots, \mathbf{p}_T]$. At the t th frame, the accumulated visual information for the activity pattern, represented as $\mathbf{P}_t = [\mathbf{p}_1, \dots, \mathbf{p}_t]$, is used for online reliable unusual activity detection. First the normalised log-likelihood of observing \mathbf{P} at the t th frame given the activity model \mathbf{M} is computed as

$$l_t = \frac{1}{t} \log P(\mathbf{P}_t|\mathbf{M}). \quad (13)$$

l_t can be easily computed online using the forward–backward procedure [41]. Specifically, to compute l_t , the K_e forward probabilities at time t are computed using the K_e forward probabilities computed at time $t-1$ together with the observations at time t (see Ref. [41] for details). Note that the complexity of computing l_t is $\mathcal{O}(K_e^2)$ and does not increase with t . We then measure how likely \mathbf{P}_t is an unusual activity using an online unusual activity measure Q_t :

$$Q_t = \begin{cases} l_1 & \text{if } t = 1, \\ (1 - \alpha)Q_{t-1} + \alpha(l_t - l_{t-1}) & \text{otherwise,} \end{cases} \quad (14)$$

where α is an accumulating factor determining how important the visual information extracted from the current frame is for unusual activity detection. We have $0 < \alpha \leq 1$. Compared to l_t as an indicator of normality/anomaly, Q_t could add more weight to more recent observations. Unusual activity is detected at frame t if

$$Q_t < Th_A, \quad (15)$$

where Th_A is the unusual activity detection threshold. The value of Th_A should be set according to the detection and false alarm rate required by each particular surveillance application. Note that it takes a time delay for Q_t to stabilise at the beginning of evaluating an activity pattern due to the nature of the forward–backward procedure. The length of this time period, denoted as T_w is related to the complexity of the MOHMM used for activity modelling. We thus set $T_w = 3K_e$ in our experiments to be reported later in Section 4.5, i.e. the unusual activity of an activity pattern is only evaluated when $t > T_w$.

At each frame t an activity pattern needs to be recognised as one of the K activity classes when it is detected as being usual,

i.e. $Q_t > Th_A$. This is achieved using an online LRT method. More specifically, we consider a hypotheses test between:

- H_k : \mathbf{P}_t is from the hypothesised model \mathbf{B}_k and belongs to the k th usual activity class,
 H_0 : \mathbf{P}_t is from a model other than \mathbf{B}_k and does not belong to the k th usual activity class,

where H_0 is called the alternative hypothesis. Using LRT, we compute the likelihood ratio of the accepting the two hypotheses as

$$r_k = \frac{P(\mathbf{P}_t; H_k)}{P(\mathbf{P}_t; H_0)}. \quad (16)$$

The hypothesis H_k can be represented by the model \mathbf{B}_k which has been learned in the activity profiling step. The key to LRT is thus to construct the alternative model which represents H_0 . In a general case, the number of possible alternatives is unlimited; $P(\mathbf{P}_t; H_0)$ can thus only be computed through approximation [42,43]. Fortunately in our case, we have determined at the t th frame that \mathbf{P}_t is usual and can only be generated by one of the K usual activity classes. Therefore, it is reasonable to construct the alternative model as a mixture of the rest $K-1$ usual activity classes. In particular, Eq. (16) is re-written as

$$r_k = \frac{P(\mathbf{P}_t|\mathbf{B}_k)}{\sum_{i \neq k} \frac{N_i}{N - N_k} P(\mathbf{P}_t|\mathbf{B}_i)}. \quad (17)$$

Note that r_k is a function of t and computed over time. \mathbf{P}_t is reliably recognised as the k th activity class only when $1 \ll Th_r < r_k$. In our experiments we found that $Th_r = 10$ led to satisfactory results. When there are more than one r_k greater than Th_r , the activity pattern is recognised as the class with the largest r_k .

For comparison, the commonly used ML method recognises \mathbf{P}_t as the k th activity class when $k = \arg \max_k \{P(\mathbf{P}_t|\mathbf{B}_k)\}$. Using the ML method, recognition has to be performed at each single frame without considering how reliable and sufficient the accumulated visual evidence is. This often causes errors especially when there are ambiguities between different classes (e.g. an activity pattern can be explained away equally well by multiple plausible activity at its early stage). Compared to the ML method, our online LRT method holds the decision on activity recognition until sufficient evidence has been accumulated to overcome ambiguities. The recognition results obtained using our approach are thus more reliable compared to those obtained by the ML method.

Table 4

Six classes of commonly occurred activity patterns in the airport scene

A1	Aircraft arrives
A2	Airbridge connected
A3	Frontal cargo service
A4	Frontal catering service
A5	Aircraft departs
A6	Airbridge disconnected

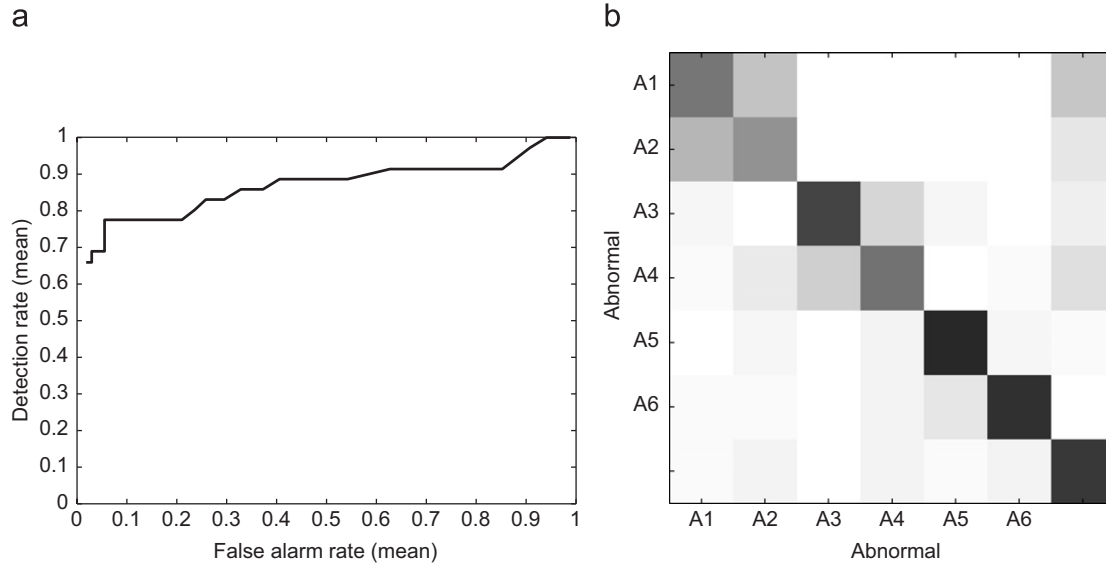


Fig. 10. The performance of unusual activity detection using aircraft docking video content model. The mean ROC curves were obtained over 20 trials. The performance of usual activity recognition obtained by averaging the results over 20 trials with $Th_A = -0.5$. Each row represents the probabilities of the corresponding class being confused with all the other classes averaged over 20 trials: (a) ROC curves; (b) confusion matrix.

4.5. Experiments

The same datasets described in Section 3.3 was employed for validating our approach on video content modelling and online unusual activity detection. The videos were segmented automatically using the FBR online segmentation algorithm described in Section 3.2, giving 59 video segments of actual activity pattern instances. Each segment has on average 428 frames with the shortest 74 and longest 2841.

Model training—A training set was created which consisted of 40 video segments and the remaining 19 were used for testing. This model training exercise was repeated 20 times and in each trial a different model was trained using a different random training set. This is in order to avoid any bias in the unusual activity detection and usual activity recognition results to be discussed later.

In each training session, the 40 training activity patterns were represented based on the event detection results and clustered to build a composite activity model. The number of clusters for each training set was determined automatically as six in every trial using the spectral clustering algorithm proposed in Ref. [20]. By observation, each of the six discovered data clusters mainly contained samples corresponding to one of the six activity classes listed in Table 4.

Unusual activity detection—Each activity pattern in a testing set was detected as being unusual when Eq. (15) was satisfied at any time after $T_w = 3K_e = 24$ frames. The accumulating factor α for computing Q_t was set to 0.1. To measure the performance of the learned models on unusual activity detection, each activity pattern in the testing sets was manually labelled as usual if there were similar patterns in the corresponding training sets and unusual otherwise. We measure the performance of unusual activity detection using unusual activity detection rate and false

Table 5

The mean and standard deviation of the unusual activity detection rate and false alarm rates

Anomaly detection rate (%)	False alarm rate (%)
79.2 ± 8.3	5.1 ± 3.9

The results were obtained over 20 trials with $Th_A = -0.5$.

alarm rate,⁵ which are defined as


$$\begin{aligned} \text{Unusual activity detection rate} &= \frac{\# \text{ True positives (unusual detected as unusual)}}{\# \text{ All positives (unusual patterns) in a dataset}} \\ \text{False alarm rate} &= \frac{\# \text{ False positives (usual detected as unusual)}}{\# \text{ All negatives (usual patterns) in a dataset}} \end{aligned} \quad (18)$$

The detection rate and false alarm rate of unusual activity detection are shown in the form of a receiver operating characteristic (ROC) curve by varying the unusual activity detection threshold Th_A (see Eq. (15)). The result is shown in Fig. 10(a) and Table 5. Figs. 11(b) and (f) show examples of online reliable unusual activity detection using our run-time unusual activity measure.


Recognition of usual activity patterns—The usual activity recognition results are illustrated in Fig. 10(b). When Th_A was set to -0.5 , the recognition rate averaged over 20 trials was

⁵ Detection rate and false alarm rate are also called true positive rate and false positive rate, respectively, in the literature. Note that the performance can also be measured using true negative rate and false negative rate. Since we have ‘true negative rate’ = $1 - \text{‘false alarm rate’}$ and ‘false negative rate’ = $1 - \text{‘detection rate’}$, showing only unusual activity detection rate and false alarm rate is adequate.

a



b



c



d




Fig. 11. Compare our LRT method with ML method for online usual activity recognition in an aircraft docking scene. (a) An unusual activity pattern where a truck brought engineers to fix a ground power cable problem. It resembled A4 in the early stage. (b) Q_t computed over time, It was detected as unusual activity from Frame 147 till the end based on Q_t . (c) r_k computed over time, The activity pattern between Frames 53–145 was recognised reliably as A4 using LRT before becoming unusual and being detected using Q_t . (d) $\log P(P_i|B_k)$ computed over time, The activity pattern was wrongly recognised as A3 between Frames 80–98 using the ML method. (e) An usual A1 activity pattern. (f) Q_t computed over time, The activity pattern was detected as usual throughout based on Q_t . (g) r_k computed over time, It was recognised reliably as A1 from Frame 73 till the end using LRT. (h) $\log P(P_i|B_k)$ computed over time, It was wrongly recognised as A2 between Frames 12–49 using ML. In (a) and (e), detected events are illustrated using the same colour scheme as in Fig. 2.

72.1%. Examples of online reliable activity recognition using our online LRT method are shown in Fig. 11 in comparison with the ML method. It can be seen that our LRT method is superior to the ML method in that usual activity patterns can be reliably and promptly recognised after sufficient visual evidence had become available to overcome the ambiguities among different activity classes.

5. Conclusion

In conclusion, we have presented a novel approach for representing, segmenting and modelling the content of CCTV surveillance videos according to the activities captured in the

scene. The video content is represented by constructing a cumulative multi-event scene vector over time. An offline MOHMM based segmentation algorithm was introduced to deal with noise in video content representation. An online FBR algorithm was also developed to detect breakpoints in the video content and segment a continuous video into activities on-the-fly. A framework for robust online activity recognition and unusual activity detection has also been proposed which is based on constructing a generative composite video content model. Experiments was carried out to evaluate the effectiveness and robustness of our approach using over 10h of challenging outdoor surveillance video footages. Our experiments demonstrate that an event based video content representation is superior to a colour

histogram based representation for surveillance video segmentation. Our experiments also suggest that an online LRT based method is preferred over the conventional ML based method for activity recognition.

It is worth pointing out that in our current MOHMM, each observational variable is 1-D and assumed to be independent to each other. Since each variable corresponds to one class of event and causal relationships exist among different classes of events, this assumption is not valid. We are investigating the possibility of joining dependent observational variables together into combined variables or introducing arcs between dependent variables to improve the performance of our MOHMM based segmentation algorithm. Similar idea has been exploited in structure improvement of the static naive Bayes classifier [44].

References

- [1] J.R. Kender, B.L. Yeo, Video scene segmentation via continuous video coherence, in: IEEE Conference on Computer Vision and Pattern Recognition, 1998, pp. 367–373.
- [2] C.W. Ngo, T.C. Pong, H.J. Zhang, Motion-based video representation for scene change detection, *Int. J. Comput. Vision* 50 (2) (1998) 127–142.
- [3] Y. Rui, S. Huang, S. Mehrota, Exploring video structures beyond the shots, in: IEEE International Conference on Multimedia Computing and Systems, 1998, pp. 237–240.
- [4] N. Babaguchi, Y. Kawai, T. Kitahashi, Event based indexing of broadcasting sports video by intermodal collaboration, *IEEE Trans. Multimedia* 4 (1) (2002) 68–75.
- [5] M.R. Naphade, I. Kozintsev, T.S. Huang, A factor graph framework for semantic indexing and retrieval in video, in: *IEEE Trans. Circuits Syst. Video Technol.*, 2002, pp. 40–52.
- [6] N. Vasconcelos, A. Lippman, Bayesian modeling of video editing and structure: semantic features for video summarization and browsing, in: *ICIP*, 1998, pp. 550–555.
- [7] G. Schwarz, Estimating the dimension of a model, *Ann. Statist.* 6 (1978) 461–464.
- [8] T. Xiang, S. Gong, D. Parkinson, Autonomous visual events detection and classification without explicit object-centred segmentation and tracking, in: *British Machine Vision Conference*, 2002, pp. 233–242.
- [9] S. Gong, T. Xiang, Recognition of group activities using dynamic probabilistic networks, in: *IEEE International Conference on Computer Vision*, 2003, pp. 742–749.
- [10] D. DeMenthon, L. Latechi, A. Rosenfeld, M. Stuckelberg, Relevance ranking of video data using hidden markov model distances and polygon simplification, in: *Advances in Visual Information Systems*, 2000.
- [11] E. Keogh, S. Chu, D. Hart, M. Pazzani, An online algorithm for segmenting time series, in: *International Conference on Data Engineering*, 2001, pp. 289–296.
- [12] L. Latecki, R. Lakamper, Convexity rule for shape decomposition based on discrete contour evolution, *Comput. Vision Image Understanding* 73 (1999) 441–454.
- [13] P. Heckbert, M. Garland, Survey of ploygonal surface simplification algorithms, in: *International Conference on Computer Graphics and Interactive Techniques*, 1997.
- [14] X. Ge, P. Smyth, Segmental semi-markov models for endpoint detection in plasma etching, in: *AEC/APC Symposium XII*, 2000.
- [15] E. Keogh, S. Kasetty, On the need for time series data mining benchmarks: a survey and empirical demonstration, *Data Min. Knowl. Discovery* 7 (4) (2003) 349–371.
- [16] H. Shatnay, S. Zdonik, Approximate queries and representations for large data sequences, in: *International Conference on Data Engineering*, 1996, pp. 546–553.
- [17] Y. Weiss, Segmentation using eigenvectors: a unifying view, in: *IEEE International Conference on Computer Vision*, 1999, pp. 975–982.
- [18] J. Shi, J. Malik, Normalized cuts and image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (8) (2000) 888–905.
- [19] S. Yu, J. Shi, Multiclass spectral clustering, in: *IEEE International Conference on Computer Vision*, 2003, pp. 313–319.
- [20] T. Xiang, S. Gong, Video behaviour profiling and abnormality detection without manual labelling, in: *IEEE International Conference on Computer Vision*, 2005, pp. 1238–1245.
- [21] O. Boiman, M. Irani, Detecting irregularities in images and in video, in: *IEEE International Conference on Computer Vision*, 2005, pp. 462–469.
- [22] N. Oliver, B. Rosario, A. Pentland, A Bayesian computer vision system for modelling human interactions, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (8) (2000) 831–843.
- [23] L. Zelnik-Manor, M. Irani, Event-based video analysis, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2001, pp. 123–130.
- [24] S. McKenna, S. Jabri, Z. Duric, A. Rosenfeld, H. Wechsler, Tracking group of people, *Comput. Vision Image Understanding* 80 (2000) 42–56.
- [25] J. Ng, S. Gong, Learning pixel-wise signal energy for understanding semantics, in: *British Machine Vision Conference*, 2001, pp. 695–704.
- [26] C. Stauffer, W. Grimson, Learning patterns of activity using real-time tracking, *IEEE Trans. Pattern Anal. Mach. Intell.* 22 (8) (2000) 747–758.
- [27] J. Sherrah, S. Gong, Automated detection of localised visual events over varying temporal scales, in: *Proceedings of the European Workshop on Advanced Video-based Surveillance System*, 2001.
- [28] A.F. Bobick, J.W. Davis, The recognition of human movement using temporal templates, *IEEE Trans. Pattern Anal. Mach. Intell.* 23 (3) (2001) 257–267.
- [29] J.H. Piater, J.L. Crowley, Multi-modal tracking of interacting targets using Gaussian approximation, in: *Proceedings of 2nd IEEE Workshop on Performance Evaluation of Tracking and Surveillance*, 2001, pp. 141–147.
- [30] T. Xiang, S. Gong, Beyond tracking: modelling activity and understanding behaviour, *Int. J. Comput. Vision* 67 (1) (2006) 21–51.
- [31] H. Greenspan, J. Goldberger, A. Mayer, Probabilistic space-time video modelling via piecewise GMM, *IEEE Trans. Pattern Anal. Mach. Intell.* 26 (3) (2004) 384–396.
- [32] H. Hung, S. Gong, Quantifying temporal saliency, in: *British Machine Vision Conference*, 2004, pp. 727–736.
- [33] C. Bishop, *Neural Networks for Pattern Recognition*, Cambridge University Press, Cambridge, 1995.
- [34] Z. Ghahramani, Learning dynamic Bayesian networks, in: *Adaptive Processing of Sequences and Data Structures. Lecture Notes in Artificial Intelligence*, Springer, Berlin, 1998, pp. 168–197.
- [35] D. Heckerman, A tutorial on learning with Bayesian networks. Technical Report MSR-TR-95-06, Microsoft Research, 1995.
- [36] C. Huang, A. Darwiche, Inference in belief networks: a procedural guide, *Int. J. Approx. Reason.* 15 (3) (1996) 225–263.
- [37] W.I. Grosky, R. Zhao, Negotiating the semantic gap: from feature maps to semantic landscapes, *Lecture Notes in Computer Science*, vol. 2234, Springer, Berlin, 2001, pp. 33–42.
- [38] P. Enser, C. Sandom, Towards a comprehensive survey of the semantic gap in visual image retrieval, in: *CIVR*, 2003, pp. 291–299.
- [39] J. Kruskal, M. Liberman, The Symmetric Time-warping Problem: From Continuous to Discrete, Addison-Wesley, 1983.
- [40] A. Dempster, N. Laird, D. Rubin, Maximum-likelihood from incomplete data via the EM algorithm, *J. R. Statist. Soc. B* 39 (1977) 1–38.
- [41] L.R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, *Proc. IEEE* 77 (2) (1989) 257–286.
- [42] J. Wilpon, L. Rabiner, C. Lee, E. Goldman, Automatic recognition of keywords in unconstrained speech using hidden markov models, in: *IEEE Trans. Acoustic Speech Signal Process.*, 1990, pp. 1870–1878.
- [43] A. Higgins, L. Bahler, J. Porter, Speaker verification using randomized phrase prompting, in: *Digital Signal Process.*, 1991, pp. 89–106.
- [44] N. Friedman, D. Geiger, M. Goldzmidt, Bayesian network classifiers, *Mach. Learn.* 29 (1997) 131–163.

About the Author—Dr. TAO XIANG is a Lecturer at the Department of Computer Science, Queen Mary, University of London. Dr. Xiang was awarded his Ph.D. in Electrical and Computer Engineering from the National University of Singapore in 2002, which involved research into 3-D Computer Vision and Visual Perception. He also received his B.Sc. degree in Electrical Engineering from Xi'an Jiaotong University in 1995, and his M.Sc. degree in Electronic Engineering from the Communication University of China (CUC) in 1998. His research interests include computer vision, image processing, statistical learning, pattern recognition, machine learning and data mining. He has been working recently on topics such as spectral clustering, video based behaviour analysis and recognition and model order selection for dynamic Bayesian networks.

About the Author—Professor SHAOGANG GONG is Professor of Visual Computation at the Department of Computer Science, Queen Mary, University of London and a Member of the UK Computing Research Committee. He heads the Queen Mary Computer Vision Group and has worked in computer vision and pattern recognition for over 20 years, published over 170 papers and a monograph. He twice won the Best Science Prize (1999 and 2001) of British Machine Vision Conferences, the Best Paper Award (2001) of IEEE International Workshops on Recognition and Tracking of Faces and Gestures and the Best Paper Award (2005) of IEE International Symposium on Imaging for Crime Detection and Prevention. He was a recipient of a Queen's Research Scientist Award (1987), a Royal Society Research Fellow (1987 and 1988), a GEC-Oxford Fellow (1989), a visiting scientist at Microsoft Research (2001), and Samsung (2003).