

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
3 October 2002 (03.10.2002)

PCT

(10) International Publication Number  
WO 02/077909 A1

- (51) International Patent Classification<sup>7</sup>: G06K 9/00, 9/34
- (21) International Application Number: PCT/US02/07185
- (22) International Filing Date: 12 March 2002 (12.03.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
09/815,385 23 March 2001 (23.03.2001) US
- (71) Applicant: DIAMONDBACK VISION, INC. [US/US];  
11600 Sunrise Valley Drive, Suite 290, Reston, VA 20191 (US).
- (72) Inventors: LIPTON, Alan, J.; 2022 Peach Orchard Drive,  
#34, Falls Church, VA 22043 (US). ALLMEN, Mark;  
6333 Starlight Drive, Morrison, CO 80465 (US). HAER-  
ING, Niels; 911 N. Barton Street, Apt. 2, Arlington, VA

22201 (US). SEVERSON, William; 8085 South Franklin Court, Littleton, CO 80122 (US). STRAT, Thomas, M.; 3227 Fox Vale Drive, Oakton, VA 22124 (US).


(74) Agent: GLUCK, Jeffrey, W.; Venable, Baetjer, Howard & Civiletti, LLP, 1201 New York Avenue, Suite 1000, P.O. Box 34385, Washington, DC 20043-9998 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: VIDEO SEGMENTATION USING STATISTICAL PIXEL MODELING



(57) Abstract: A method for segmenting video data into foreground and background portions utilizes statistical modeling (316) of the pixels. A statistical model of the background (312) is built for each pixel, and each pixel in an incoming video frame is compared (314) with the background statistical model for that pixel. Pixels are determined to be foreground or background based on the comparison.

WO 02/077909 A1



European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

**Published:**

— *with international search report*

**VIDEO SEGMENTATION USING STATISTICAL PIXEL MODELING*****Field of the Invention***

5           The present invention relates to processing of video frames for use in video processing systems, for example, video compression systems. More specifically, it deals with ways of segmenting video frames into their component parts using statistical properties of regions comprising the video frames.

***Background of the Invention***

10           In object-based video compression, as well as in other types of object-oriented video processing, the input video is separated into two streams. One stream contains the information representing stationary background information, and the other stream contains information representing the moving portions of the video, to be denoted as  
15 foreground information. The background information is represented as a background model, including a scene model, i.e., a composite image composed from a series of related images, as, for example, one would find in a sequence of video frames; the background model may also contain additional models and modeling information. Scene models are generated by aligning images (for example, by matching points and/or  
20 regions) and determining overlap among them; generation of scene models is discussed in further depth in commonly-assigned U.S. Patent Applications Serial Nos. 09/472,162, filed December 27, 1999, and 09/609,919, filed July 3, 2000, both incorporated by reference in their entireties herein. In an efficient transmission or storage scheme, the scene model need be transmitted only once, while the foreground information is  
25 transmitted for each frame. For example, in the case of an observer (i.e., camera or the like, which is the source of the video) that undergoes only pan, tilt, roll, and zoom types of motion, the scene model need be transmitted only once because the appearance of the scene model does not change from frame to frame, except in a well-defined way based on the observer motion, which can be easily accounted for by transmitting motion  
30 parameters. Note that such techniques are also applicable in the case of other forms of motion, besides pan, tilt, roll, and zoom.

To make automatic object-oriented video processing feasible, it is necessary to be able to distinguish the regions in the video sequence that are moving or changing and to

separate (i.e., segment) them from the stationary background regions. This segmentation must be performed in the presence of apparent motion, for example, as would be induced by a panning, tilting, rolling, and/or zooming observer (or due to other motion-related phenomena, including actual observer motion). To account for this motion, images are first aligned; that is, corresponding locations in the images (i.e., frames) are determined, as discussed above. After this alignment, objects that are truly moving or changing, relative to the stationary background, can be segmented from the stationary objects in the scene. The stationary regions are then used to create (or to update) the scene model, and the moving foreground objects are identified for each frame.

10           It is not an easy thing to identify and automatically distinguish between video objects that are moving foreground and stationary background, particularly in the presence of observer motion, as discussed above. Furthermore, to provide the maximum degree of compression or the maximum fineness or accuracy of other video processing techniques, it is desirable to segment foreground objects as finely as possible; this  
15           enables, for example, the maintenance of smoothness between successive video frames and crispness within individual frames. Known techniques have proven, however, to be difficult to utilize and inaccurate for small foreground objects and have required excessive processing power and memory. It would, therefore, be desirable to have a technique that permits accurate segmentation between the foreground and background  
20           information and accurate, crisp representations of the foreground objects, without the limitations of prior techniques.

### *Summary of the Invention*

          The present invention is directed to a method for segmentation of video into  
25           foreground information and background information, based on statistical properties of the source video. More particularly, the method is based on creating and updating statistical information pertaining to a characteristic of regions of the video and the labeling of those regions (i.e., as foreground or background) based on the statistical information. For example, in one embodiment, the regions are pixels, and the characteristic is chromatic  
30           intensity. Many other possibilities exist, as will become apparent.

          In embodiments of the invention, a background model is developed containing at least two components. A first component is the scene model, which may be built and

updated, for example, as discussed in the aforementioned U.S. patent applications. A second component is a background statistical model.

In a first embodiment, the inventive method comprises a two-pass process of video segmentation. The two passes of the embodiment comprise a first pass in which a background statistical model is built and updated and a second pass in which regions in the frames are segmented. An embodiment of the first pass comprises steps of aligning each video frame with a scene model and updating the background statistical model based on the aligned frame data. An embodiment of the second pass comprises, for each frame, steps of labeling regions of the frame and performing spatial filtering.

In a second embodiment, the inventive method comprises a one-pass process of video segmentation. The single pass comprises, for each frame in a frame sequence of a video stream, steps of aligning the frame with a scene model; building a background statistical model; labeling the regions of the frame, and performing spatial/temporal filtering.

In yet another embodiment, the inventive method comprises a modified version of the aforementioned one-pass process of video segmentation. This embodiment is similar to the previous embodiment, except that the step of building a background statistical model is replaced with a step of building a background statistical model *and* a secondary statistical model.

Each of these embodiments may be embodied in the forms of a computer system running software executing their steps and a computer-readable medium containing software representing their steps.

### *Definitions*

In describing the invention, the following definitions are applicable throughout (including above).

A "computer" refers to any apparatus that is capable of accepting a structured input, processing the structured input according to prescribed rules, and producing results of the processing as output. Examples of a computer include: a computer; a general purpose computer; a supercomputer; a mainframe; a super mini-computer; a mini-computer; a workstation; a micro-computer; a server; an interactive television; a hybrid combination of a computer and an interactive television; and application-specific hardware to emulate a computer and/or software. A computer can have a single processor

or multiple processors, which can operate in parallel and/or not in parallel. A computer also refers to two or more computers connected together via a network for transmitting or receiving information between the computers. An example of such a computer includes a distributed computer system for processing information via computers linked by a network.

A “computer-readable medium” refers to any storage device used for storing data accessible by a computer. Examples of a computer-readable medium include: a magnetic hard disk; a floppy disk; an optical disk, like a CD-ROM or a DVD; a magnetic tape; a memory chip; and a carrier wave used to carry computer-readable electronic data, such as those used in transmitting and receiving e-mail or in accessing a network.

“Software” refers to prescribed rules to operate a computer. Examples of software include: software; code segments; instructions; computer programs; and programmed logic.

A “computer system” refers to a system having a computer, where the computer comprises a computer-readable medium embodying software to operate the computer.

A “network” refers to a number of computers and associated devices that are connected by communication facilities. A network involves permanent connections such as cables or temporary connections such as those made through telephone or other communication links. Examples of a network include: an internet, such as the Internet; an intranet; a local area network (LAN); a wide area network (WAN); and a combination of networks, such as an internet and an intranet.

“Video” refers to motion pictures represented in analog and/or digital form. Examples of video include television, movies, image sequences from a camera or other observer, and computer-generated image sequences. These can be obtained from, for example, a live feed, a storage device, a firewire interface, a video digitizer, a computer graphics engine, or a network connection.

“Video processing” refers to any manipulation of video, including, for example, compression and editing.

A “frame” refers to a particular image or other discrete unit within a video.

### ***Brief Description of the Drawings***

The invention will now be described in further detail in connection with the attached drawings, in which:

Figure 1 shows a flowchart corresponding to an implementation of a first embodiment of the invention;

Figures 2a and 2b show flowcharts corresponding to two alternative embodiments of the labeling step in the flowchart of Figure 1;

5 Figure 3a and 3b show flowcharts corresponding to implementations of the spatial/temporal filtering step in the flowchart of Figure 1;

Figure 4 shows a flowchart corresponding to an implementation of a second embodiment of the invention;

10 Figure 5 shows a flowchart corresponding to an implementation of one of the steps in the flowchart of Figure 4;

Figures 6a and 6b together show a flowchart corresponding to an implementation of another one of the steps in the flowchart of Figure 4;

Figure 7 shows a flowchart corresponding to an implementation of a third embodiment of the invention;

15 Figures 8a and 8b together show a flowchart corresponding to an implementation of one of the steps in the flowchart of Figure 7; and

Figure 9 depicts an embodiment of the invention in the form of software embodied on a computer-readable medium, which may be part of a computer system.

20 Note that identical objects are labeled with the same reference numerals in all of the drawings that contain them.

### ***Detailed Description of the Invention***

As discussed above, the present invention is directed to the segmentation of video streams into foreground information, which corresponds to moving objects, and  
25 background information, which corresponds to the stationary portions of the video. The present invention may be embodied in a number of ways, of which three specific ones are discussed below. These embodiments are meant to be exemplary, rather than exclusive.

The ensuing discussion refers to “pixels” and “chromatic intensity;” however, the inventive method is not so limited. Rather, the processing may involve any type of region  
30 (including regions comprising multiple pixels), not just a pixel, and may use any type of characteristic measured with respect to or related to such a region, not just chromatic intensity.

#### **1. First Embodiment -- Two-Pass Segmentation**

The first embodiment of the invention is depicted in Figure 1 and corresponds to a two-pass method of segmentation. As shown in Figure 1, the method begins by obtaining a frame (or video) sequence from a video stream (Step 1). The frame sequence preferably includes two or more frames of the video stream. The frame sequence can be, for  
5 example, a portion of the video stream or the entire video stream. As a portion of the video stream, the frame sequence can be, for example, one continuous sequence of frames of the video stream or two or more discontinuous sequences of frames of the video stream. As part of the alignment step, the scene model is also built and updated.

After Step 1, in Step 2, it is determined whether or not all frames have yet been  
10 processed. If not, the next frame is taken and aligned with the underlying scene model of the video stream (Step 3); such alignment is discussed above, and more detailed discussions of alignment techniques may be found, for example, in commonly-assigned U.S. Patent Applications Serial Nos. 09/472,162, filed December 27, 1999, and 09/609,919, filed July 3, 2000, both incorporated by reference in their entireties herein, as  
15 discussed above, as well as in numerous other references.

The inventive method is based on the use of statistical modeling to determine whether a particular pixel should be classified as being a foreground object or a part thereof or as being the background or a part thereof. Step 4 deals with the building and updating of a statistical model of the background, using each frame aligned in Step 3.

20 The statistical model of the present invention comprises first- and second-order statistics. In the ensuing discussion, mean and standard deviation will be used as such first- and second-order statistics; however, this is meant to be merely exemplary of the statistics that may be used.

In general, the mean of  $N$  samples,  $\bar{x}$ , is computed by taking the sum of the  
25 samples and dividing it by  $N$ , i.e.,

$$\bar{x} = \frac{\sum_{i=1}^N x_i}{N}, \quad (1)$$

where  $x_i$  is a particular sample corresponding to a given pixel (or region), which in the present case could be, for example, the measured chromatic intensity of the  $i^{\text{th}}$  sample corresponding to the given pixel (or region). In the present setting, then, such a mean  
30 would be computed for each pixel or region.



While Eqn. (1) gives the general formula for a sample mean, it may not always be optimal to use this formula. In video processing applications, a pixel's sample value may change drastically when an object moves through the pixel and change (drastically) back to a value around its previous value after the moving object is no longer within that pixel.

5 In order to address this type of consideration, the invention utilizes a weighted average, in which the prior values are weighted more heavily than the present value. In particular, the following equation may be used:

$$\bar{x}_N = W_p \bar{x}_{N-1} + W_n x_N, \quad (2)$$

where  $W_p$  is the weight of the past values and  $W_n$  is the weight assigned to the newest value. Additionally,  $\bar{x}_J$  represents the weighted average taken over  $J$  samples, and  $x_K$  represents the  $K^{\text{th}}$  sample.  $W_p$  and  $W_n$  may be set to any pair of values between zero and one such that their sum is one and such that  $W_n < W_p$ , so as to guarantee that the past values are more heavily weighted than the newest value. As an example, the inventors have successfully used  $W_p=0.9$  and  $W_n=0.1$ .

15 Standard deviation,  $\sigma$ , is determined as the square root of the variance,  $\sigma^2$ , of the values under consideration. In general, variance is determined by the following formula:

$$\sigma^2 = \overline{x^2} - (\bar{x})^2, \quad (3)$$

where  $\overline{x^2}$  represents the average of  $x^2$ ; thus, the standard deviation is given by

$$\sigma = \sqrt{\overline{x^2} - (\bar{x})^2}. \quad (4)$$

20 Because the inventive method uses running statistics, this becomes

$$\sigma_N = \sqrt{\left\{ \overline{x^2} \right\}_N - (\bar{x}_N)^2}, \quad (4a)$$

where  $\bar{x}_N$  is as defined in Eqn. (2) above, and  $\left\{ \overline{x^2} \right\}_N$  is defined as the weighted average of the squared values of the samples, through the  $N^{\text{th}}$  sample, and is given by

$$\left\{ \overline{x^2} \right\}_N = W_p \left\{ \overline{x^2} \right\}_{N-1} + W_n x_N^2. \quad (5)$$

25 As in the case of the weighted average of the sample values, the weights are used to assure that past values are more heavily weighted than the present value.

Given this, Step 4 works to create and update the statistical model by computing the value of Eqn. (4a) for each pixel, for each frame. In Step 4, the values for the pixels are also stored on a pixel-by-pixel basis (as opposed to how they are received, i.e., on a frame-by-frame basis); that is, an array of values is compiled for each pixel over the

sequence of frames. Note that in an alternative embodiment, Step 4 only performs this storage of values.

Following Step 4, the method returns to Step 2 to check whether or not all of the frames have been processed. If they have, then the method proceeds to Step 5, which  
5 commences the second pass of the embodiment.

In Step 5, the statistical background model is finalized. This is done by using the stored values for each pixel and determining their mode, the mode being the value that occurs most often. This may be accomplished, for example, by taking a histogram of the stored values and selecting the value for which the histogram has the highest value. The  
10 mode of each pixel is then assigned as the value of the background statistical model for that pixel.

Following Step 5, the method proceeds to Step 6, which determines whether or not all of the frames have been processed yet. If not, then the method proceeds to Step 7, in which each pixel in the frame is labeled as being a foreground (FG) pixel or a  
15 background (BG) pixel. Two alternative embodiments of the workings of this step are shown in the flowcharts of Figures 2a and 2b.

Figure 2a depicts a two decision level method. In Figure 2a, the pixel labeling Step 7 begins with Step 71, where it is determined whether or not all of the pixels in the frame have been processed. If not, then the method proceeds to Step 72 to examine the  
20 next pixel. Step 72 determines whether or not the pixel matches the background statistical model, i.e., whether the value of the pixel matches the mode for that pixel. This is performed by taking the absolute difference between the pixel value and the value of the background statistical model for the pixel (i.e., the mode) and comparing it with a threshold; that is,

$$25 \quad \Delta = |x_{pixel} - m_{pixel}| \quad (6)$$

is compared with a threshold  $\theta$ . In Eqn. (6),  $x_{pixel}$  denotes the value of the pixel, while  $m_{pixel}$  represents the value of the statistical background model for that pixel.

The threshold  $\theta$  may be determined in many ways. For example, it may be taken to be a function of standard deviation (of the given pixel),  $\sigma$ . In a particular exemplary  
30 embodiment,  $\theta = 3\sigma$ ; in another embodiment,  $\theta = K\sigma$ , where  $K$  is chosen by the user. As another example,  $\theta$  may be assigned a predetermined value (again, for each pixel) or one chosen by the user.

If  $\Delta \leq \theta$ , then the pixel value is considered to match the background statistical model. In this case, the pixel is labeled as background (BG) in Step 73, and the algorithm proceeds back to Step 71. Otherwise, if  $\Delta > \theta$ , then the pixel value is considered not to match the background statistical model, and the pixel is labeled as foreground (FG) in Step 74. Again, the algorithm then proceeds back to Step 71. If Step 71 determines that all of the pixels (in the frame) have been processed, then Step 7 is finished.

Figure 2b depicts a three decision level method, labeled 7'. In Figure 2b, the process once again begins with Step 71, a step of determining whether or not all pixels have yet been processed. If not, the process considers the next pixel to be processed and executes Step 72, the step of determining whether or not the pixel being processed matches the background statistical model; this is done in the same way as in Figure 2a. If yes, then the pixel is labeled as BG (Step 73), and the process loops back to Step 71. If not, then the process proceeds to Step 75; this is where the process of Figure 2b is distinguished from that of Figure 2a.

In Step 75, the process determines whether or not the pixel under consideration is far from matching the background statistical model. This is accomplished via a threshold test similar to Step 72, only in Step 75,  $\theta$  is given a larger value. As in Step 72,  $\theta$  may be user-assigned or predetermined. In one embodiment,  $\theta = N\sigma$ , where  $N$  is either a predetermined or user-set number,  $N > K$ . In another embodiment,  $N=6$ .

If the result of Step 75 is that  $\Delta \leq \theta$ , then the pixel is labeled as FG (Step 74). If not, then the pixel is labeled definite foreground (DFG), in Step 76. In each case, the process loops back to Step 71. Once Step 71 determines that all pixels in the frame have been processed, Step 7' is complete.

Returning to Figure 1, once all of the pixels of a frame have been labeled, the process proceeds to Step 8, in which spatial/temporal filtering is performed. While shown as a sequential step in Figure 1, Step 8 may alternatively be performed in parallel with Step 7. Details of Step 8 are shown in the flowcharts of Figures 3a and 3b.

In Figure 3a, Step 8 commences with a test as to whether or not all the pixels of the frame have been processed (Step 81). If not, in Step 85, the algorithm selects the next pixel,  $P_i$ , for processing and proceeds to Step 82, where it is determined whether or not the pixel is labeled as BG. If it is, then the process goes back to Step 81. If not, then the pixel undergoes further processing in Steps 83 and 84.

Step 83, neighborhood filtering, is used to correct for misalignments when the images are aligned. If the current image is slightly misaligned with the growing background statistical model, then, particularly near strong edges, the inventive segmentation procedure, using the background statistical model, will label pixels as foreground. Neighborhood filtering will correct for this. An embodiment of Step 83 is depicted in the flowchart of Figure 3b.

In Figure 3b, Step 83 begins with Step 831, where a determination is made of the scene model location,  $P_m$ , corresponding to  $P_i$ . Next, a neighborhood, comprising the pixels,  $P'_m$ , surrounding  $P_m$  in the scene model, is selected (Step 832). Step 833 next determines if all of the pixels in the neighborhood have been processed. If yes, Step 83 is complete, and the label of  $P_i$  remains as it was; if not, the process proceeds to Step 834, where the next neighborhood pixel  $P'_m$  is considered. Step 835 then tests to determine whether or not  $P_i$  matches  $P'_m$ . This matching test is accomplished by executing the labeling step (Step 7 or 7') in a modified fashion, using  $P_i$  as the pixel under consideration and  $P'_m$  as the "corresponding" background statistical model point. If the labeling step returns a label of FG or DFG, there is no match, whereas if it returns a label of BG, there is a match. If there is no match, the process loops back to Step 833; if there is a match, then this is an indication that  $P_i$  might be mislabeled, and the process continues to Step 836. In Step 836, a neighborhood, comprising the pixels,  $P'_i$ , surrounding  $P_i$  in the frame, is selected, and an analogous process is performed. That is, in Step 833, it is determined whether or not all of the pixels,  $P'_i$  in the neighborhood have yet been considered. If yes, then Step 83 is complete, and the label of  $P_i$  remains as it was; if not, then the process proceeds to Step 838, where the next neighborhood pixel,  $P'_i$ , is considered. Step 839 tests to determine if  $P_m$  matches  $P'_i$ ; this is performed analogously to Step 833, with the  $P'_i$  under consideration being used as the pixel being considered and  $P_m$  as its "corresponding" background statistical model point. If it does not, then the process loops back to Step 837; if it does, then  $P_i$  is relabeled as BG, and Step 83 is complete.

Returning to Figure 3a, following Step 83, Step 84 is executed, in which morphological erosions and dilations are performed. First, a predetermined number,  $n$ , of erosions are performed to remove incorrectly labeled foreground. Note that pixels labeled DFG may not be eroded because they represent either a pixel that is almost certainly foreground. This is followed by  $n$  dilations, which restore the pixels that were

correctly labeled as foreground but were eroded. Finally, a second predetermined number,  $m$ , of dilations are performed to fill in holes in foreground objects. The erosions and dilations may be performed using conventional erosion and dilation techniques, applied in accordance with user-specified parameters, and modified, as discussed above, such that pixels labeled DFG are not eroded.

In alternative embodiments, Step 84 may comprise filtering techniques other than or in addition to morphological erosions and dilations. In general, Step 84 may employ any form or forms of spatial and/or temporal filtering.

Returning to Figure 1, following Step 8, the algorithm returns to Step 6, to determine whether or not all frames have been processed. If yes, then the processing of the frame sequence is complete, and the process ends (Step 9).

This two-pass embodiment has the advantage of relative simplicity, and it is an acceptable approach for applications not requiring immediate or low-latency processing. Examples of such applications include off-line video compression and non-linear video editing. On the other hand, many other applications do have such requirements, and the embodiments to be discussed below are tailored to address these requirements.

## **2. Second Embodiment - One-Pass Segmentation**

Figure 4 depicts a flowchart of a one-pass segmentation process, according to a second embodiment of the invention. Comparing Figure 4 with Figure 1 (the first embodiment), the second embodiment differs in that there is only a single pass of processing for each frame sequence. This single pass, as shown in Steps 2, 3, 31, 32, 8 in Figure 4, incorporates the processes of the second pass (Steps 5-8 in Figure 1) with the first pass (Steps 2-4 in Figure 1), albeit in a modified form, as will be discussed below.

As in the case of the first embodiment, the second embodiment (one-pass process), shown in Figure 4, begins by obtaining a frame sequence (Step 1). As in the first embodiment, the process then performs a test to determine whether or not all of the frames have yet been processed (Step 2). Also as in the first embodiment, if the answer is no, then the next frame to be processed is aligned with the scene model (Step 3). As discussed above, the scene model component of the background model is built and updated as part of Step 3, so there is always at least a deterministically-determined value in the background model at each location.

At this point, the process includes a step of building a background statistical model (Step 31). This differs from Step 4 of Figure 1, and is depicted in further detail in

Figure 5. The process begins with a step of determining whether or not all pixels in the frame being processed have been processed (Step 311). If not, then the process determines whether or not the background statistical model is “mature” (Step 312) and “stable” (Step 313).

5           The reason for Steps 312 and 313 is that, initially, the statistical background model will not be sufficiently developed to make accurate decisions as to the nature of pixels. To overcome this, some number of frames should be processed before pixels are labeled (i.e., the background statistical model should be “mature”); in one embodiment of the present invention, this is a user-defined parameter. This may be implemented as a  
10 “look-ahead” procedure, in which a limited number of frames are used to accumulate the background statistical model prior to pixel labeling (Step 32 in Figure 4).

          While simply processing a user-defined number of frames may suffice to provide a mature statistical model, stability is a second concern (Step 313), and it depends upon the standard deviation of the background statistical model. In particular, as will be  
15 discussed below, the statistical background model includes a standard deviation for each pixel. The statistical model (for a particular pixel) is defined as having become “stable” when its variance (or, equivalently, its standard deviation) is reasonably small. In an embodiment of the present invention, Step 313 determines this by comparing the standard deviation with a user-defined threshold parameter; if the standard deviation is less than  
20 this threshold, then the statistical background model (for that pixel) is determined to be stable.

          As to the flow of Step 31, in Figure 5, if the background statistical model is determined to be mature (Step 312), it is determined whether or not the background statistical model is stable (Step 313). If either of these tests (Steps 312 and 313) fails, the  
25 process proceeds to Step 315, in which the background statistical model of the pixel being processed is updated using the current value of that pixel. Step 315 will be explained further below.

          If the background statistical model is determined to be both mature and stable (in Steps 312 and 313), the process proceeds to Step 314, where it is determined whether or  
30 not the pixel being processed matches the background statistical model. If yes, then the background statistical model is updated using the current pixel value (Step 315); if no, then the process loops back to Step 311 to determine if all pixels in the frame have been processed.

Step 314 operates by determining whether or not the current pixel value is within some range of the mean value of the pixel, according to the current background statistical model. In one embodiment of the invention, the range is a user-defined range. In yet another embodiment, it is determined to be a user-defined number of standard deviations; i.e., the pixel value,  $x$ , matches the background statistical model if

$$\left| x_{pixel} - \overline{x_{pixel}} \right| \leq K\sigma, \quad (7)$$

where  $K$  is the user-defined number of standard deviations,  $\sigma$ ;  $x_{pixel}$  is the current pixel value; and  $\overline{x_{pixel}}$  is the mean value of the current pixel in the background statistical

model. The purpose of performing Step 314 is to ensure, to the extent possible, that only background pixels are used to develop and update the background statistical model.

In Step 315, the background statistical model is updated. In this embodiment, the background statistical model consists of the mean and standard deviation of the values for each pixel (over the sequence of frames). These are computed according to Eqns. (2) and (4a) above.

Following Step 315, the process loops back to Step 311, to determine if all pixels (in the current frame) have been processed. Once all of the pixels have been processed, the process proceeds to Step 316, where the background statistical model is finalized. This finalization consists of assigning to each pixel its current mean value and standard deviation (i.e., the result of processing all of the frames up to that point).

Note that it is possible for the background statistical model for a given pixel never to stabilize. This generally indicates that the particular pixel is not a background pixel in the sequence of frames, and there is, therefore, no need to assign it a value for the purposes of the background statistical model. Noting that, as discussed above, a scene model is also built and updated, there is always at least a deterministically-determined value associated with each pixel in the background model.

Following Step 316, the process goes to Step 32, as shown in Figure 4, where the pixels in the frame are labeled according to their type (i.e., definite foreground, foreground or background). Step 32 is shown in further detail in the flowchart of Figures 6a and 6b.

The following concepts are embodied in the description of Step 32 to follow. Ideally, labeling would always be done by testing each pixel against its corresponding

point in the background statistical model, but this is not always possible. If the background statistical model is not ready to use on the basis of number of frames processed (i.e., “mature”), then the process must fall back on testing against the corresponding point in the scene model. If the background statistical model is ready to use but has not yet settled down (i.e., is not “stable”), this is a sign that the pixel is varying and should be labeled as being foreground. If the background statistical model has, for some reason (i.e., because it fails to match the scene model or because it has become unsettled again), become unusable, the process must once again fall back on testing against the scene model.

10           As shown in Figure 6a, Step 32 begins with Step 321, where it is determined whether or not all pixels (in the current frame) have been processed. If yes, Step 32 is complete; if not, the next pixel is processed in Steps 322 *et seq.*

          Step 322 determines whether or not the background statistical model is mature. This is done in the same manner as in Step 312 of Figure 5, discussed above. If not, the process proceeds to Step 323, where it is determined whether or not the pixel matches the background chromatic data of the corresponding point of the scene model.

15           Step 323 is performed by carrying out a test to determine whether or not the given pixel falls within some range of the background chromatic data value. This is analogous to Step 314 of Figure 5, substituting the background chromatic data value for the statistical mean. The threshold may be determined in a similar fashion (predetermined, user-determined, or the like).

          If Step 323 determines that the pixel does match the background chromatic data, then the pixel is labeled BG (following connector A) in Step 329 of Figure 6b. From Step 329, the process loops back (via connector D) to Step 321.

25           If Step 323 determines that the pixel does not match the background chromatic data, then the pixel is labeled FG (following connector B) in Step 3210 of Figure 6b. From the Step 3210, the process loops back (via connector D) to Step 321.

          If Step 322 determines that the background statistical model is mature, processing proceeds to Step 324, which determines whether or not the background statistical model is stable. Step 324 performs this task in the same manner as Step 313 of Figure 5, discussed above. If not, the process proceeds to Step 325, where it is determined if the background statistical model was ever stable (i.e., if it was once stable but is now unstable). If yes, then the process branches to Step 323, and the process proceeds from



there as described above. If no, the pixel is labeled DFG (following connector C) in Step 3211 of Figure 6b, after which the process loops back (via connector D) to Step 321.

If Step 324 determines that the background statistical model is stable, the process goes to Step 326. Step 326 tests whether the background statistical model matches the background chromatic data. Similar to the previous matching tests above, this test takes an absolute difference between the value of the background statistical model (i.e., the mean) for the pixel and the background chromatic data (i.e., of the scene model) for the pixel. This absolute difference is then compared to some threshold value, as above (predetermined, user-determined, or the like).

If Step 326 determines that there is not a match between the background statistical model and the background chromatic data, the process branches to Step 323, where processing proceeds in the same fashion as described above. If Step 326, on the other hand, determines that there is a match, the process continues to Step 327.

Step 327 determines whether or not the current pixel matches the background statistical model. This step is performed in the same manner as Step 314 of Figure 5, discussed above. If the current pixel does match (which, as discussed above, is determined by comparing it to the mean value corresponding to the current pixel), the pixel is labeled BG (following connector A) in Step 329 of Figure 6b, and the process then loops back (via connector D) to Step 321. If not, then further testing is performed in Step 328.

Step 328 determines whether, given that the current pixel value does not reflect a BG pixel, it reflects a FG pixel or a DFG pixel. This is done by determining if the pixel value is far from matching the background statistical model. As discussed above, a FG pixel is distinguished from a BG pixel (in Step 325) by determining if its value differs from the mean by more than a particular amount, for example, a number of standard deviations (see Eqn. (7)). Step 328 applies the same test, but using a larger range. Again, the threshold may set as a predetermined parameter, as a computed parameter, or as a user-defined parameter, and it may be given in terms of a number of standard deviations from the mean, i.e.,

$$|x_{pixel} - \overline{x_{pixel}}| \leq N\sigma, \quad (8)$$

where  $N$  is a number greater than  $K$  of Eqn. (7). If the pixel value lies outside the range defined, for example, by Eqn. (8), it is labeled DFG (following connector C) in Step 3211

of Figure 6b, and the process loops back (via connector D) to Step 321. If it lies within the range, the pixel is labeled FG (following connector B) in Step 3210 of Figure 6b, and the process proceeds (via connector D) to Step 321.

After Step 32 is complete, the process proceeds to Step 8, as shown in Figure 4, where spatial/temporal filtering is performed on the pixels in the frame. Step 8 is implemented, in this embodiment of the invention, in the same manner in which it is implemented for the two-pass embodiment, except that the pixel labeling algorithm of Figures 6a and 6b is used for Steps 833 and 837 of Step 83 (as opposed to the pixel labeling algorithms used in the two-pass embodiment). Following Step 8, the process loops back to Step 2, where, if all frames have been processed, the process ends.

A single-pass approach, like the one present here, has the advantage of not requiring a second pass, thus, reducing the latency associated with the process. This is useful for applications in which high latencies would be detrimental, for example, video teleconferencing, webcasting, real-time gaming, and the like.

### 15 **3. Third Embodiment - Modified One-Pass Segmentation**

While the one-pass approach described above has a lower latency than the two-pass approach, it does have a disadvantage in regard to the background statistical model. In particular, the cumulative statistical modeling approach used in the one-pass embodiment of the invention may stabilize on a non-representative statistical model for an element (i.e., pixel, region, etc.; that is, whatever size element is under consideration). If the values (e.g., chromatic values) of frame elements corresponding to a particular element of the video scene fundamentally change (i.e., something happens to change the video, for example, a parked car driving away, a moving car parking, the lighting changes, etc.), then the scene model element will no longer accurately represent the true scene. This can be addressed by utilizing a mechanism for dynamically updating the background statistical model so that at any given time it accurately represents the true nature of the scene depicted in the video. Such a mechanism is depicted in the embodiment of the invention shown in Figure 7.

In Figure 7, Steps 1-3, 32, 8, and 9 are as described in the one-pass embodiment above. The embodiment of Figure 7 differs from that of Figure 4 in that after a given frame is aligned with the scene model (Step 3), the process executes Step 310, in which the background statistical model and, simultaneously, a secondary background statistical model are built. Step 310 is more fully described in connection with Figures 8a and 8b.

As shown in Figure 8a, Step 310 includes all of the steps shown in Step 31 in Figure 5 (which are shown using the same reference numerals), and it begins with a step of determining whether or not all pixels have yet been processed (Step 311). If not, the next pixel is processed by proceeding to Step 312. In Step 312, it is determined whether or not the background statistical model is mature. If not, the process branches to Step 315, where the pixel is used to update the background statistical model. Following Step 315, the process loops back to Step 311.

If Step 312 determines that the background statistical model is mature, the process proceeds to Step 313, where it is determined whether or not the background statistical model is stable. If it is not, then, as in the case of a negative determination in Step 312, the process branches to Step 315 (and then loops back to Step 311). Otherwise, the process proceeds to Step 314.

In Step 314, it is determined whether or not the pixel under consideration matches the background statistical model. If it does, the process proceeds with Step 315 (and then loops back to Step 311); otherwise, the process executes the steps shown in Figure 8b, which build and update a secondary background statistical model. This secondary background statistical model is built in parallel with the background statistical model, as reflected in Figure 8b; uses the same procedures as are used to build and update the background statistical model; and represents the pixel values that do not match the background statistical model.

Following a negative determination in Step 314, the process then makes a determination as to whether or not the secondary background statistical model is mature (Step 3107). This determination is made in the same fashion as in Step 313. If not, the process branches to Step 3109, where the secondary background statistical model is updated, using the same procedures as for the background statistical model (Step 315). From Step 3109, the process loops back to Step 311 (in Figure 8a).

If Step 3107 determines that the secondary background statistical model is mature, the process proceeds to Step 3108, which determines (using the same procedures as in Step 314) whether or not the secondary background statistical model is stable. If not, the process proceeds to Step 3109 (and from there to Step 311). If yes, then the process branches to Step 31010, in which the background statistical model is replaced with the secondary background statistical model, after which the process loops back to Step 311. Additionally, concurrently with the replacement of the background statistical model by

the secondary background statistical model in Step 31010, the scene model data is replaced with the mean value of the secondary statistical model. At this point, the secondary background statistical model is reset to zero, and a new one will be built using subsequent data.

5           This modified one-pass embodiment has the advantage of improved statistical accuracy over the one-pass embodiment, and it solves the potential problem of changing background images. It does this while still maintaining improved latency time over the two-pass embodiment, and at only a negligible decrease in processing speed compared with the one-pass embodiment.

#### 10   **4. Additional Embodiments and Remarks**

          While the above discussion considers two-level and three-level pixel labeling algorithms, this embodiment is not limited only to these cases. Indeed, it is contemplated that an arbitrary number of decision levels, corresponding to different ranges (i.e., threshold values) may be used. In such a case, a form of soft-decision logic could be used  
15 to make decisions in subsequent steps of the segmentation process.

          The above discussion primarily discusses pixels and chromatic values (which may be RGB, YUV, intensity, etc.); however, as discussed above, the invention is not limited to these quantities. Regions other than pixels may be used, and quantities other than chromatic values may be used.

20           As discussed above, the invention, including all of the embodiments discussed in the preceding sections, may be embodied in the form of a computer system or in the form of a computer-readable medium containing software implementing the invention. This is depicted in Figure 9, which shows a plan view for a computer system for the invention. The computer 91 includes a computer-readable medium 92 embodying software for  
25 implementing the invention and/or software to operate the computer 91 in accordance with the invention. Computer 91 receives a video stream and outputs segmented video, as shown. Alternatively, the segmented video may be further processed within the computer.

          The invention has been described in detail with respect to preferred embodiments, and it will now be apparent from the foregoing to those skilled in the art that changes and  
30 modifications may be made without departing from the invention in its broader aspects. The invention, therefore, as defined in the appended claims, is intended to cover all such changes and modifications as fall within the true spirit of the invention.

### Claims

We claim:

1. A two-pass method of video segmentation, for differentiating between foreground and background portions of video, comprising the steps of:
  - 5 obtaining a frame sequence from an input video stream;
  - executing a first-pass method for each frame of the frame sequence, the first-pass method comprising the steps of:
    - aligning the frame with a scene model; and
    - updating a background statistical model; and
    - 10 finalizing the background statistical model;
  - executing a second-pass method for each frame of the frame sequence, the second-pass method comprising the steps of:
    - labeling each region of the frame; and
    - performing spatial/temporal filtering of the regions of the frame.
- 15 2. A computer system comprising:
  - a computer; and
  - a computer-readable medium containing software implementing the method of Claim 1.
3. A computer-readable medium comprising software implementing the method of  
20 Claim 1.
4. The method according to Claim 1, wherein the step of updating a background statistical model comprises the steps of:
  - compiling values corresponding to the regions of the frames; and
  - computing a standard deviation for each region of the frames.
- 25 5. The method according to Claim 1, wherein the step of finalizing the background statistical model comprises the step of:
  - for each region of the frames, determining a statistical mode based on the compiled values; and
  - assigning the statistical mode of each region as its value in the background  
30 statistical model.
6. The method according to Claim 1, wherein the step of labeling each region of the frame comprises the steps of:

determining whether or not the region matches the background statistical model;  
and

assigning a label based on the result of the determining step.

7. The method according to Claim 6, wherein the step of determining comprises the  
5 step of determining whether a value of the region of the frame falls within some range  
about a corresponding value of the region in the background statistical model; and  
wherein the step of assigning a label comprises the steps of:

assigning a first label if the value of the region of the frame does fall  
within the range; and

10 assigning a second label if the value of the region of the frame does not  
fall within the range.

8. The method according to Claim 7, wherein said first label is background and said  
second label is foreground.

9. The method according to Claim 6, wherein the step of determining comprises the  
15 steps of:

determining whether a value of the region of the frame falls within a first range  
about a corresponding value of the region in the background statistical model; and  
if the value of the region of the frame does not fall within the first range,  
determining if it falls within a second range about said corresponding value of the region  
20 in the background statistical model; and

wherein the step of assigning a label comprises the steps of:

assigning a first label if the value of the region of the frame falls within the  
first range;

25 assigning a second label if the value of the region of the frame does not  
fall within the first range and does lie within the second range; and

assigning a third label if the value of the region of the frame does not fall  
within the second range.

10. The method according to Claim 9, wherein said first label is background, said  
second label is foreground, and said third label is definite foreground.

30 11. The method according to Claim 1, wherein the step of labeling each region of the  
frame comprises the steps of:

determining whether or not the region matches the background statistical model,  
comprising the step of performing determinations as to whether or not a value of the

region of the frame falls within each of an arbitrary number of ranges about a corresponding value of the region in the background statistical model ; and

*assigning a label based on the result of the determining step.*

12. The method according to Claim 1, wherein the step of spatial/temporal filtering  
5 comprises the steps of, for each region of the frame:

performing neighborhood filtering; and

performing spatial filtering, temporal filtering, or a combination thereof.

13. The method according to Claim 12, wherein the step of performing spatial  
filtering, temporal filtering, or a combination thereof comprises the step of performing

10 morphological erosions and dilations, wherein the morphological erosions and dilations are performed in two dimensions or in three dimensions.

14. The method according to Claim 13, wherein said step of performing  
morphological erosions and dilations comprises the steps of:

15 performing a number,  $n$ , of erosions, wherein any region labeled definite foreground is not eroded;

following the step of performing  $n$  erosions, performing  $n$  dilations; and

following the step of performing  $n$  dilations, performing a second number,  $m$ , of dilations.

15. The method according to Claim 12, wherein the step of performing neighborhood  
20 filtering comprises the steps of, for each region of the frame,  $P_i$ , not labeled background:

determining a location,  $P_m$ , in the scene model corresponding to  $P_i$ ;

determining a neighborhood of regions,  $P'_m$  around  $P_m$ ;

determining if any of the  $P'_m$  matches  $P_i$ ;

if any of the  $P'_m$  matches  $P_i$ , determining a neighborhood of regions  $P'_i$  around  $P_i$ ;

25 determining if any of the  $P'_i$  matches  $P_m$ ; and

if any of the  $P'_m$  matches  $P_i$ , relabeling  $P_i$  background.

16. A one-pass method of video segmentation, for differentiating between foreground  
and background portions of video, comprising the steps of:

obtaining a frame sequence from a video stream; and

30 for each frame in the frame sequence, performing the following steps:

aligning the frame with a scene model;

building a background statistical model;

labeling the regions of the frame; and

performing spatial/temporal filtering.

17. A computer system comprising:

a computer; and

a computer-readable medium containing software implementing the method of

5 Claim 16.

18. A computer-readable medium comprising software implementing the method of Claim 16.

19. The method according to Claim 16, wherein the step of spatial/temporal filtering comprises the steps of, for each region of the frame:

10 performing neighborhood filtering; and

performing spatial filtering, temporal filtering, or a combination thereof.

20. The method according to Claim 19, wherein the step of performing spatial filtering, temporal filtering, or a combination thereof comprises the step of performing morphological erosions and dilations, wherein the morphological erosions and dilations  
15 are performed in two dimensions or in three dimensions.

21. The method according to Claim 20, wherein said step of performing morphological erosions and dilations comprises the steps of:

performing a number,  $n$ , of erosions, wherein any region labeled definite foreground is not eroded;

20 following the step of performing  $n$  erosions, performing  $n$  dilations; and

following the step of performing  $n$  dilations, performing a second number,  $m$ , of dilations.

22. The method according to Claim 19, wherein the step of performing neighborhood filtering comprises the steps of, for each region of the frame,  $P_i$ , not labeled background:

25 determining a location,  $P_m$ , in the scene model corresponding to  $P_i$ ;

determining a neighborhood of regions,  $P'_m$  around  $P_m$ ;

determining if any of the  $P'_m$  matches  $P_i$ ;

if any of the  $P'_m$  matches  $P_i$ , determining a neighborhood of regions  $P'_i$  around  $P_i$ ;

determining if any of the  $P'_i$  matches  $P_m$ ; and

30 if any of the  $P'_m$  matches  $P_i$ , relabeling  $P_i$  background.

23. The method according to Claim 16, wherein said step of building a background statistical model comprises the steps of:

for each region of the frame, performing the steps of:



determining if the background statistical model is mature and stable;

if the background statistical model is mature and stable, determining if a value of the region matches a corresponding value of the background statistical model; and

5 if the value of the region matches the value of the background statistical model, updating the background statistical model using the value of the region.

24. The method according to Claim 23, wherein the step of determining if the background statistical model is mature and stable comprises the steps of:

10 determining if the background statistical model is mature, and if it is not, updating the background statistical model using the value of the region;

if the background statistical model is determined to be mature, determining if the background statistical model is stable;

if the background statistical model is determined not to be stable, updating the background statistical model using the value of the region.

15 25. The method according to Claim 24, wherein each step of updating the background statistical model using the value of the region comprises the steps of:

updating a running mean value for the region; and

updating a running standard deviation value for the region.

20 26. The method according to Claim 23, wherein the step of updating the background statistical model using the value of the region comprises the steps of:

updating a running mean value for the region; and

updating a running standard deviation value for the region.

27. The method according to Claim 23, wherein said step of building a background statistical model further comprises the step of finalizing the background statistical model.

25 28. The method according to Claim 16, wherein said step of labeling the regions of the frame comprises, for each region of the frame, the steps of:

making determinations as to the maturity and stability of the background statistical model for that region;

30 determining if a value of the background statistical model for that region matches a value of the scene model for that region;

determining if a value of the region matches the value of the scene model for that region; and

determining to what extent the value of the region matches the value of the background statistical model for that region;

and wherein, based on the results of the making and determining steps, a label is assigned to the region.

5 29. The method according to Claim 28, wherein the step of making determinations as to the maturity and stability of the background statistical model comprises the steps of:

determining whether or not the background statistical model is mature;

determining whether or not the background statistical model is stable; and

if the background statistical model is not stable, determining if it was ever

10 previously stable.

30. The method according to Claim 29, wherein the step of determining if the background statistical model is stable is executed only if the background statistical model is determined to be mature; and wherein, if the background statistical model is determined not to be mature, the step of determining if a value of the region matches the value of the scene model for that region is executed.

15

31. The method according to Claim 29, wherein, if the background statistical model is determined to be stable, the step of determining if a value of the background statistical model for that region matches a value of the scene model for that region is executed.

32. The method according to Claim 29, wherein, if the step of determining if the background statistical model was ever previously stable returns an affirmative result, the step of determining if a value of the region matches the value of the scene model for that region is executed; and wherein, if the step of determining if the background statistical model was ever previously stable returns a negative result, then the region is labeled definite foreground.

20

25 33. The method according to Claim 28, wherein if the step of determining if a value of the region matches the value of the scene model for that region returns an affirmative result, the region is labeled background, and if it returns a negative result, the region is labeled foreground.

25

34. The method according to Claim 28, wherein, if the step of determining if a value of the background statistical model for that region matches a value of the scene model for that region returns an affirmative result, the step of determining to what extent the value of the region matches the value of the background statistical model for that region is executed; and wherein, if the step of determining if a value of the background statistical

30

model for that region matches a value of the scene model for that region returns a negative result, then the step of determining if a value of the region matches the value of the scene model for that region is executed.

35. The method according to Claim 28, wherein the step of determining to what extent the value of the region matches the value of the background statistical model for that region comprises the steps of:

determining if the value of the region matches the value of the background statistical model for that region;

10 if the step of determining if the value of the region matches the value of the background statistical model for that region, labeling the region background; and

if the step of determining if the value of the region matches the value of the background statistical model for that region returns a negative result, determining if the value of the region falls far from the background statistical model.

36. The method according to Claim 35, wherein the step of determining if the value of the region falls far from the background statistical model comprises the step of performing a threshold comparison; and wherein:

if the step of determining if the value of the region falls far from the background statistical model returns an affirmative result, labeling the region definite foreground, and

20 if the step of determining if the value of the region falls far from the background statistical model returns a negative result, labeling the region foreground.

37. A one-pass method of video segmentation, for differentiating between foreground and background portions of video, comprising the steps of:

obtaining a frame sequence from a video stream; and

for each frame in the frame sequence, performing the following steps:

25 aligning the frame with a scene model;

building a background statistical model and a secondary statistical model;

labeling the regions of the frame; and

performing spatial/temporal filtering.

38. A computer system comprising:

30 a computer; and

a computer-readable medium containing software implementing the method of

Claim 37.

39. A computer-readable medium comprising software implementing the method of Claim 37.

40. The method according to Claim 37, wherein the step of spatial/temporal filtering comprises the steps of, for each region of the frame:

5 performing neighborhood filtering; and  
performing spatial filtering, temporal filtering, or a combination thereof.

41. The method according to Claim 40, wherein the step of performing spatial filtering, temporal filtering, or a combination thereof comprises the step of performing morphological erosions and dilations, wherein the morphological erosions and dilations  
10 are performed in two dimensions or in three dimensions.

42. The method according to Claim 41, wherein said step of performing morphological erosions and dilations comprises the steps of:

performing a number,  $n$ , of erosions, wherein any region labeled definite foreground is not eroded;  
15 following the step of performing  $n$  erosions, performing  $n$  dilations; and  
following the step of performing  $n$  dilations, performing a second number,  $m$ , of dilations.

43. The method according to Claim 40, wherein the step of performing neighborhood filtering comprises the steps of, for each region of the frame,  $P_i$ , not labeled background:

20 determining a location,  $P_m$ , in the scene model corresponding to  $P_i$ ;  
determining a neighborhood of regions,  $P'_m$  around  $P_m$ ;  
determining if any of the  $P'_m$  matches  $P_i$ ;  
if any of the  $P'_m$  matches  $P_i$ , determining a neighborhood of regions  $P'_i$  around  $P_i$ ;  
determining if any of the  $P'_i$  matches  $P_m$ ; and  
25 if any of the  $P'_m$  matches  $P_i$ , relabeling  $P_i$  background.

44. The method according to Claim 37, wherein said step of building a background statistical model comprises the steps of:

for each region of the frame, performing the steps of:  
determining if the background statistical model is mature and stable;  
30 if the background statistical model is mature and stable, determining if a value of the region matches a corresponding value of the background statistical model;

if the value of the region matches the value of the background statistical model, updating the background statistical model using the value of the region;  
and

if the value of the region does not match the value of the background statistical model, updating the secondary statistical model using the value of the region.

5  
45. The method according to Claim 44, wherein the step of determining if the background statistical model is mature and stable comprises the steps of:

10 determining if the background statistical model is mature, and if it is not, updating the background statistical model using the value of the region;

if the background statistical model is determined to be mature, determining if the background statistical model is stable;

if the background statistical model is determined not to be stable, updating the background statistical model using the value of the region.

15 46. The method according to Claim 45, wherein each step of updating the background statistical model using the value of the region comprises the steps of:

updating a running mean value for the region; and

updating a running standard deviation value for the region.

20 47. The method according to Claim 44, wherein the step of updating the background statistical model using the value of the region comprises the steps of:

updating a running mean value for the region; and

updating a running standard deviation value for the region.

48. The method according to Claim 44, wherein said step of updating the secondary statistical model comprises the steps of:

25 determining if the secondary statistical model is mature;

if the secondary statistical model is not mature, updating the secondary statistical model using the value of the region;

if the secondary statistical model is mature, determining if the secondary statistical model is stable;

30 if the secondary statistical model is not stable, updating the secondary statistical model using the value of the region; and

if the secondary statistical model is stable, replacing the background statistical model for the region with the secondary statistical model for the region.

49. The method according to Claim 44, wherein said step of building a background statistical model further comprises the step of finalizing the background statistical model.

50. The method according to Claim 37, wherein said step of labeling the regions of the frame comprises, for each region of the frame, the steps of:

5 making determinations as to the maturity and stability of the background statistical model for that region;

determining if the value of a background statistical model for that region matches a value of the scene model for that region;

10 determining if a value of the region matches the value of the scene model for that region; and

determining to what extent the value of the region matches the value of the background statistical model for that region;

and wherein, based on the results of the making and determining steps, a label is assigned to the region.

15 51. The method according to Claim 50, wherein the step of making determinations as to the maturity and stability of the background statistical model comprises the steps of:

determining whether or not the background statistical model is mature;

determining whether or not the background statistical model is stable; and

20 if the background statistical model is not stable, determining if it was ever previously stable.

52. The method according to Claim 51, wherein the step of determining if the background statistical model is stable is executed only if the background statistical model is determined to be mature; and wherein, if the background statistical model is determined not to be mature, the step of determining if a value of the region matches the value of the scene model for that region is executed.

53. The method according to Claim 51, wherein, if the background statistical model is determined to be stable, the step of determining if a value of the background statistical model for that region matches a value of the scene model for that region is executed.

54. The method according to Claim 51, wherein, if the step of determining if the background statistical model was ever previously stable returns an affirmative result, the step of determining if a value of the region matches the value of the scene model for that region is executed; and wherein, if the step of determining if the background statistical

model was ever previously stable returns a negative result, then the region is labeled definite foreground.

55. The method according to Claim 50, wherein if the step of determining if a value of the region matches the value of the scene model for that region returns an affirmative result, the region is labeled background, and if it returns a negative result, the region is labeled foreground.

56. The method according to Claim 50, wherein, if the step of determining if a value of the background statistical model for that region matches a value of the scene model for that region returns an affirmative result, the step of determining to what extent the value of the region matches the value of the background statistical model for that region is executed; and wherein, if the step of determining if a value of the background statistical model for that region matches a value of the scene model for that region returns a negative result, then the step of determining if a value of the region matches the value of the scene model for that region is executed.

57. The method according to Claim 50, wherein the step of determining to what extent the value of the region matches the value of the background statistical model for that region comprises the steps of:

determining if the value of the region matches the value of the background statistical model for that region;

if the step of determining if the value of the region matches the value of the background statistical model for that region, labeling the region background; and

if the step of determining if the value of the region matches the value of the background statistical model for that region returns a negative result, determining if the value of the region falls far from the background statistical model.

58. The method according to Claim 57, wherein the step of determining if the value of the region falls far from the background statistical model comprises the step of performing a threshold comparison; and wherein:

if the step of determining if the value of the region falls far from the background statistical model returns an affirmative result, labeling the region definite foreground, and

if the step of determining if the value of the region falls far from the background statistical model returns a negative result, labeling the region foreground.

59. A method of segmenting video into foreground and background components, comprising the steps of:

obtaining a sequence of video frames;  
building and updating at least one background statistical model for each region of  
the video frames, based on the video frames; and  
assigning labels to the regions, based on the at least one background statistical  
5 model.

60. The method according to Claim 59, further comprising the step of:  
filtering the results of the step of assigning.

61. The method according to Claim 60, wherein the step of filtering comprises the  
steps of:

10 performing neighborhood filtering; and  
performing spatial filtering, temporal filtering, or a combination thereof.

62. The method according to Claim 59, wherein the step of building and updating at  
least one background statistical model comprises the steps of:

15 building and updating a primary background statistical model; and  
building and updating a secondary background statistical model;

wherein, if the primary background statistical model is mature and stable, if data  
corresponding to a given region does not match the primary background statistical model,  
and if the secondary background statistical model is mature and stable, then the secondary  
background statistical model is used instead of the primary background statistical model.



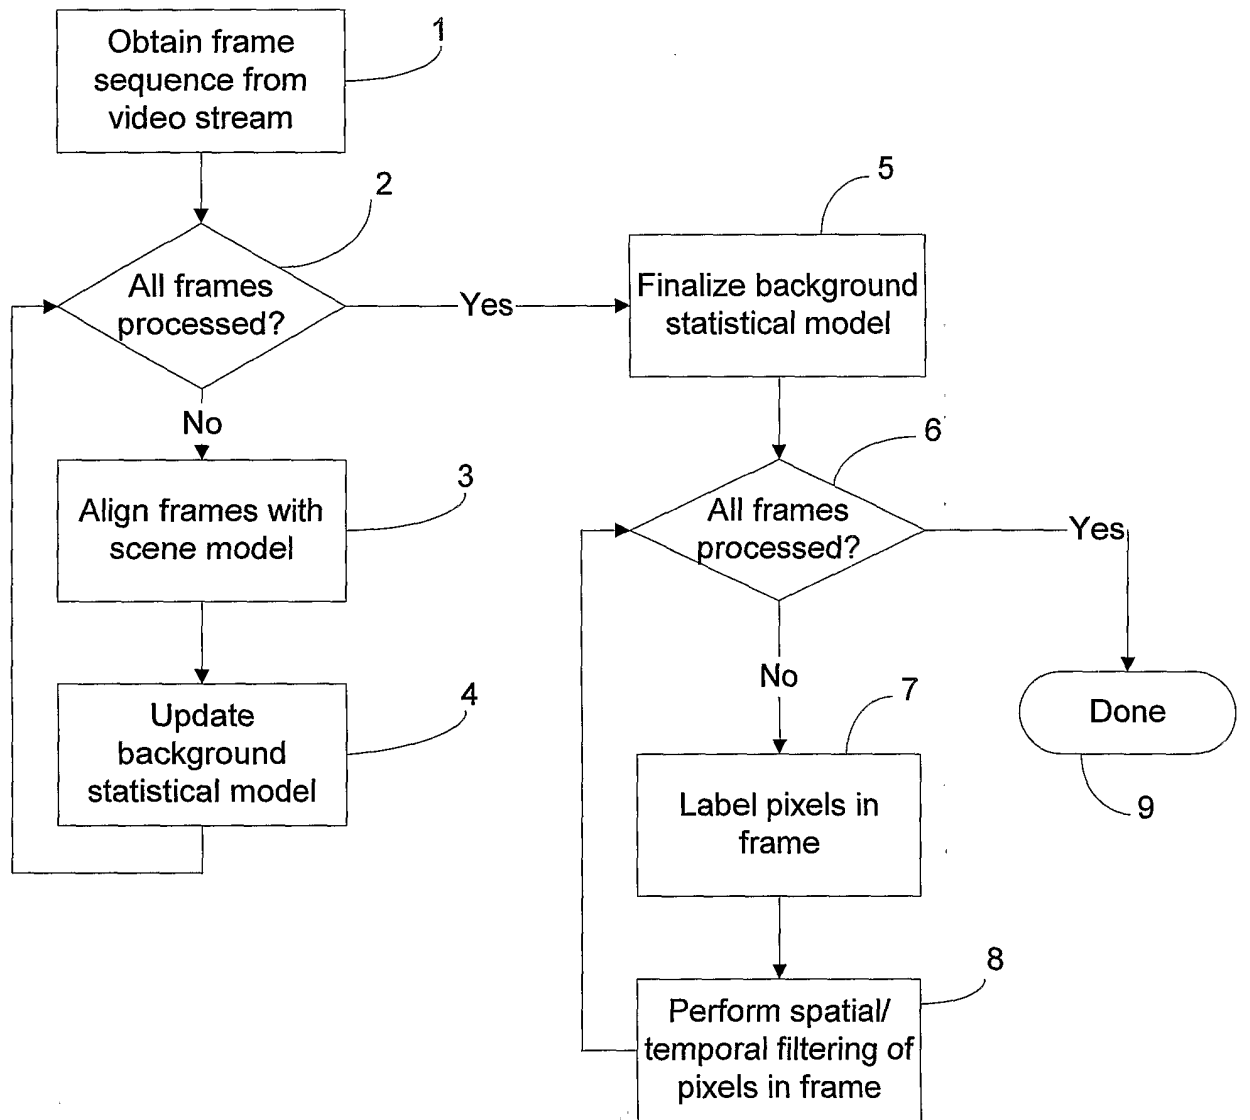


Figure 1




Figure 2a

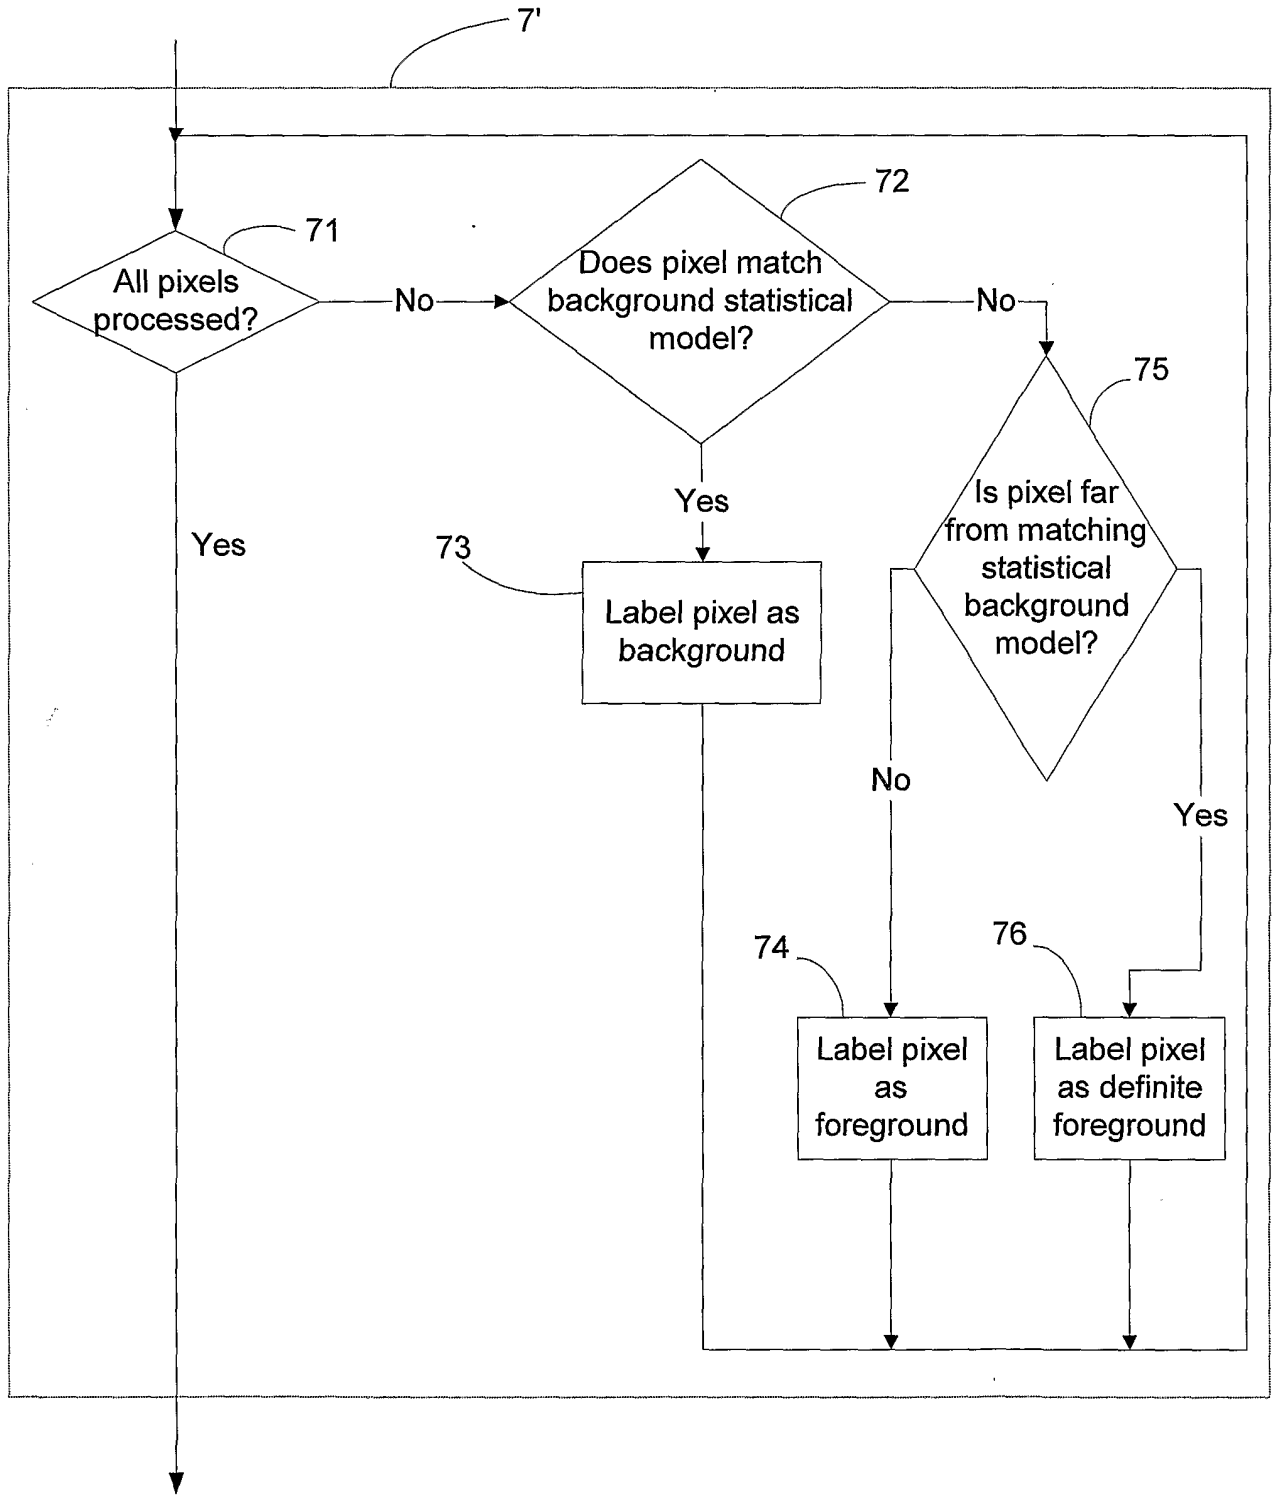


Figure 2b

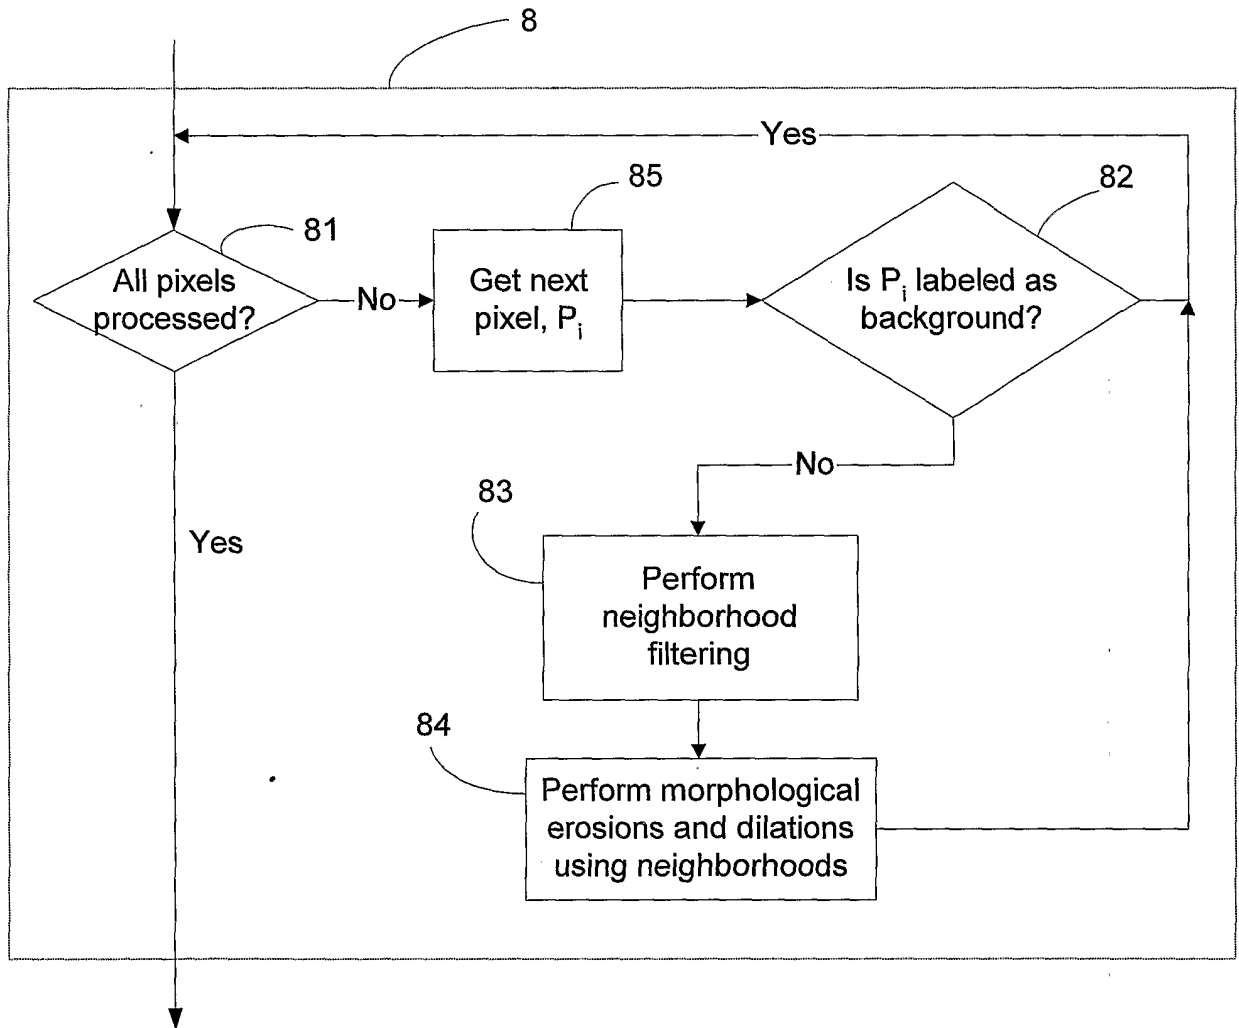


Figure 3a

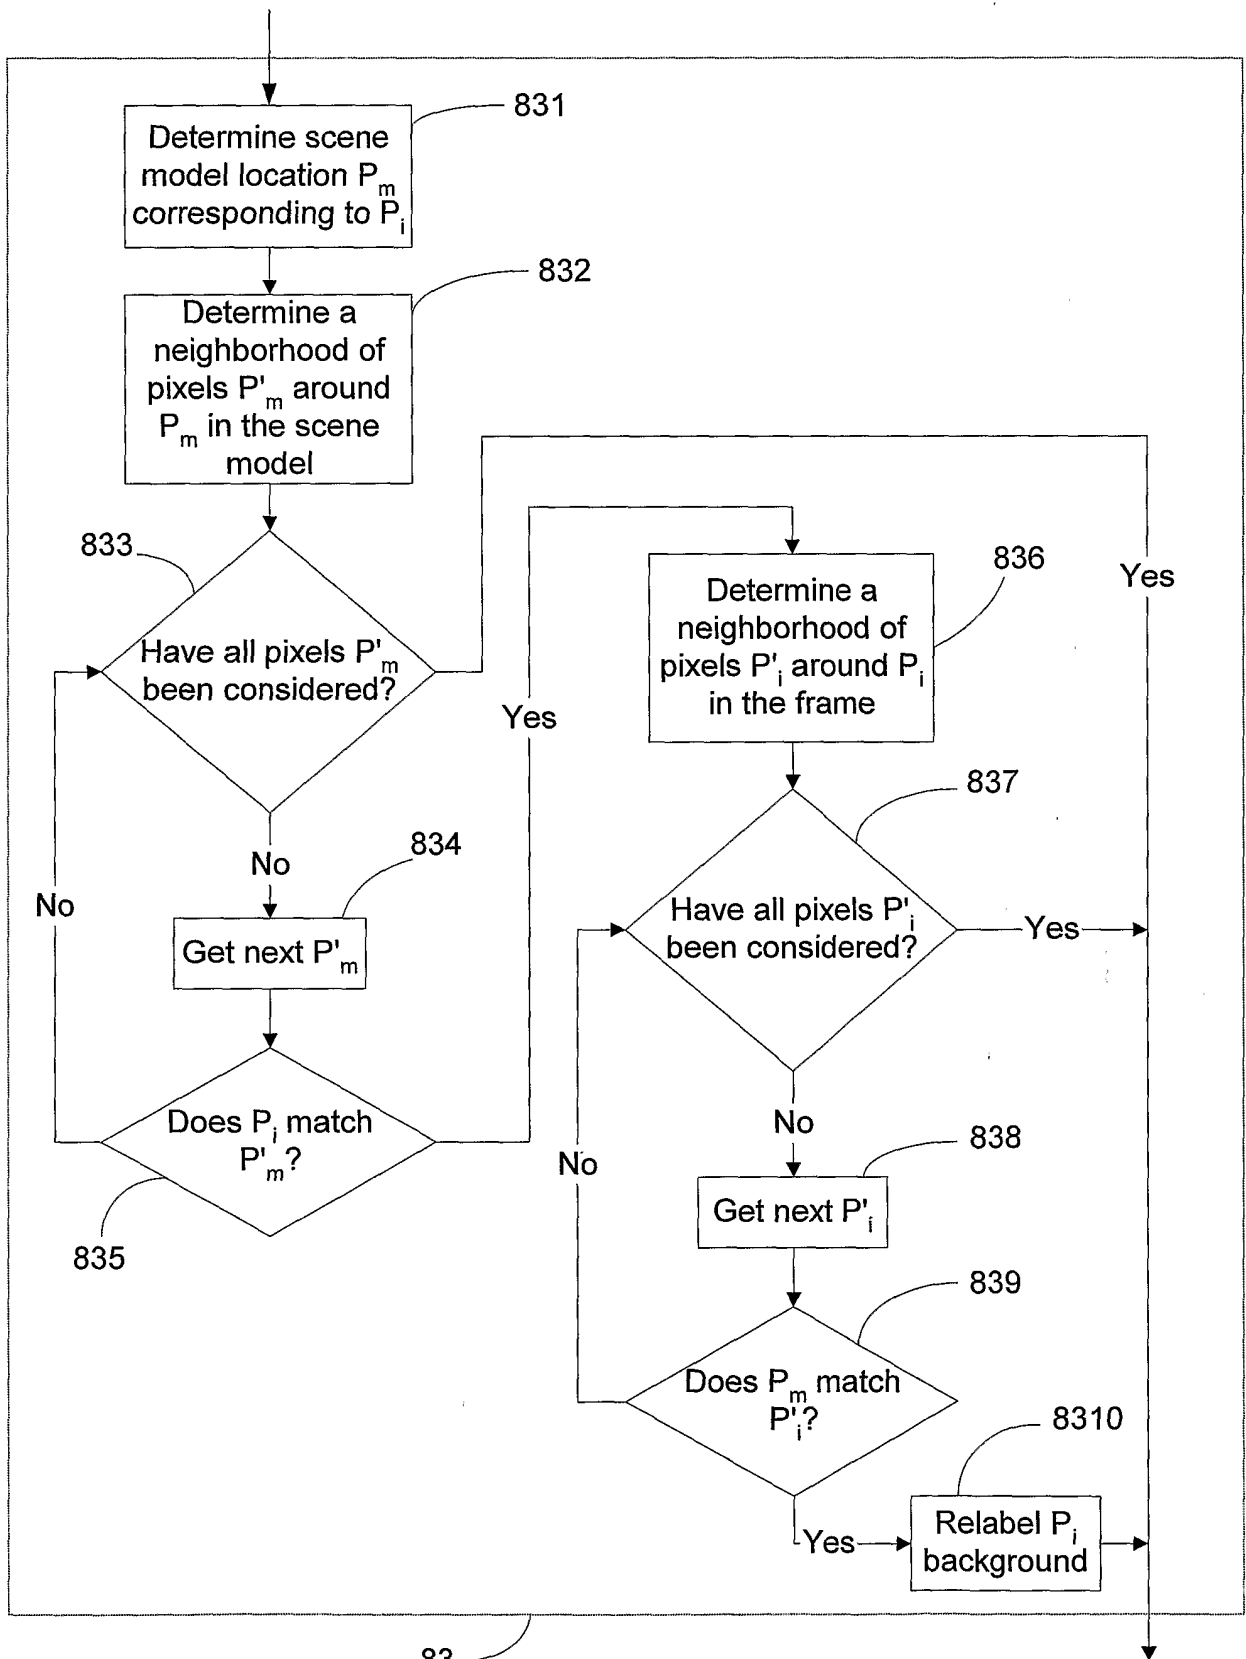


Figure 3b

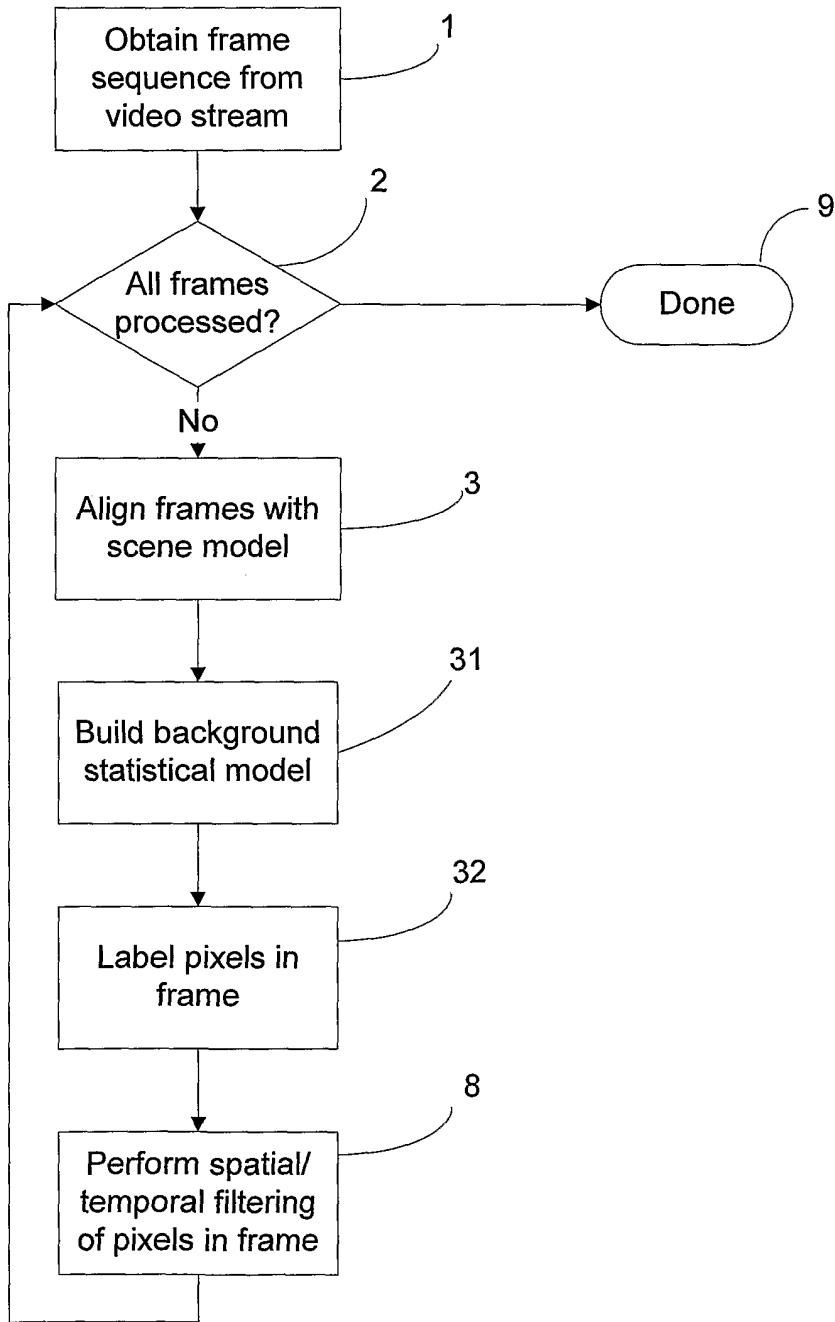


Figure 4




Figure 5




Figure 6a






Figure 6b




Figure 7




Figure 8a

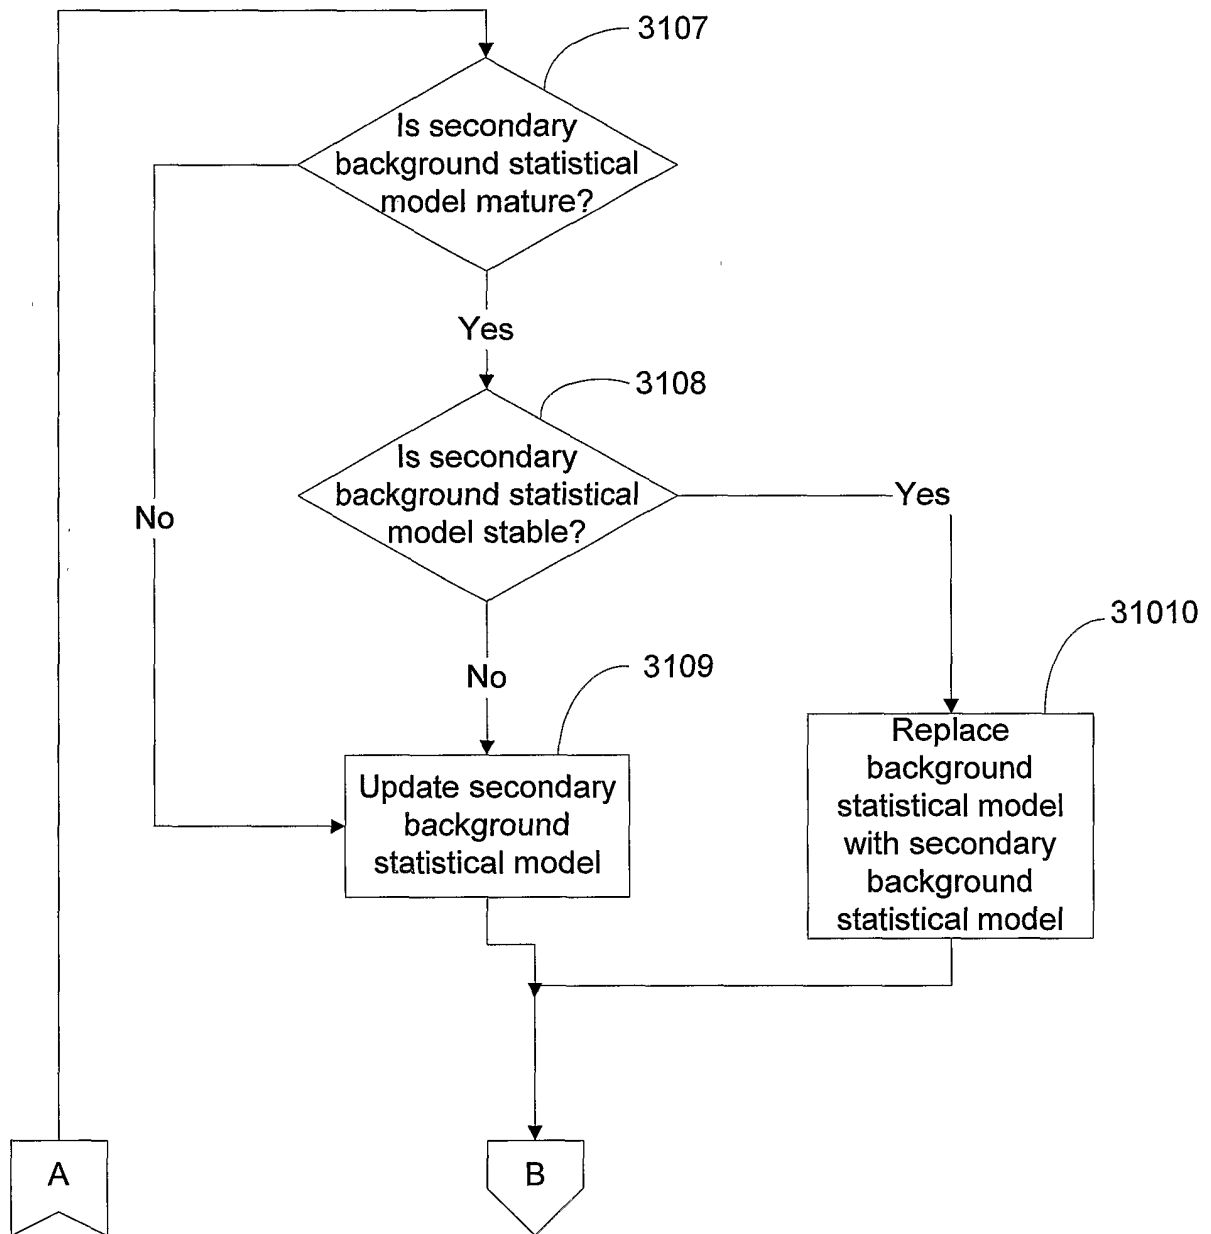


Figure 8b




Figure 9

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US02/07185

| <p><b>A. CLASSIFICATION OF SUBJECT MATTER</b><br/>                 IPC(7) : G06K 9/00; 9/34<br/>                 US CL : 348/586,588; 375/240.01,240.29; 382/173,180,190,203,210<br/>                 According to International Patent Classification (IPC) or to both national classification and IPC</p>  |   |   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
|--|---|---|---|---|-----------------------|---|---|---------------------------|---|---|------|---|---|-----------------------|
| <p><b>B. FIELDS SEARCHED</b></p> <p>Minimum documentation searched (classification system followed by classification symbols)<br/>                 U.S. : 348/586,588; 375/240.01,240.29; 382/173,180,190,203,210</p> <p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched</p> <p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)</p>   |   |   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| <p><b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b></p> <table border="1"> <thead> <tr> <th>Category *</th> <th>Citation of document, with indication, where appropriate, of the relevant passages</th> <th>Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td>X</td> <td>US 5,764,306 A (STEFFANO) 09 June 1998, column 1, lines 53-64</td> <td>1-4,16-18,37-39 and 59-60</td> </tr> <tr> <td>Y</td> <td>US 4,764,971 A (SULLIVAN) 16 August 1988, column 1, lines 12-60</td> <td>5-15</td> </tr> <tr> <td>Y</td> <td>US 5,990,955 A (KOZ) 23 November 1999, column 25, lines 50-67 and column 26, lines 1-40</td> <td>19-36,40-58 and 61-62</td> </tr> </tbody> </table>  |   |   | Category *  | Citation of document, with indication, where appropriate, of the relevant passages  | Relevant to claim No. | X | US 5,764,306 A (STEFFANO) 09 June 1998, column 1, lines 53-64 | 1-4,16-18,37-39 and 59-60 | Y | US 4,764,971 A (SULLIVAN) 16 August 1988, column 1, lines 12-60 | 5-15 | Y | US 5,990,955 A (KOZ) 23 November 1999, column 25, lines 50-67 and column 26, lines 1-40 | 19-36,40-58 and 61-62 |
| Category *   | Citation of document, with indication, where appropriate, of the relevant passages  | Relevant to claim No.   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| X  | US 5,764,306 A (STEFFANO) 09 June 1998, column 1, lines 53-64   | 1-4,16-18,37-39 and 59-60   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| Y  | US 4,764,971 A (SULLIVAN) 16 August 1988, column 1, lines 12-60   | 5-15  |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| Y  | US 5,990,955 A (KOZ) 23 November 1999, column 25, lines 50-67 and column 26, lines 1-40   | 19-36,40-58 and 61-62   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| <p><input type="checkbox"/> Further documents are listed in the continuation of Box C.      <input type="checkbox"/> See patent family annex.</p>  |   |   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| <p>* Special categories of cited documents:</p> <table border="0"> <tr> <td style="vertical-align: top;"> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"B" earlier application or patent published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </td> <td style="vertical-align: top;"> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&amp;" document member of the same patent family</p> </td> </tr> </table> |   |   | <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"B" earlier application or patent published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> | <p>"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</p> <p>"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</p> <p>"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</p> <p>"&amp;" document member of the same patent family</p> |                       |   |   |                           |   |   |      |   |   |                       |
| <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"B" earlier application or patent published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>  | <p>"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</p> <p>"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</p> <p>"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</p> <p>"&amp;" document member of the same patent family</p> |   |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| <p>Date of the actual completion of the international search</p> <p>23 May 2002 (23.05.2002)</p>   |   | <p>Date of mailing of the international search report</p> <p>08 JUL 2002</p>  |   |   |                       |   |   |                           |   |   |      |   |   |                       |
| <p>Name and mailing address of the ISA/US<br/>                 Commissioner of Patents and Trademarks<br/>                 Box PCT<br/>                 Washington, D.C. 20231<br/>                 Facsimile No. (703)305-3230</p>  |   | <p>Authorized officer<br/>                 PHUOC TRAN<br/>                 Telephone No. 703-305-4861</p> <p><i>Rugenia Zagan</i></p> |   |   |                       |   |   |                           |   |   |      |   |   |                       |