

ROBUST CAMERA PAN VECTOR ESTIMATION USING ITERATIVE CENTER OF **MASS**

BACKGROUND

[0001] This invention relates to improving video and graphics processing.

[0002] Standard video and film content for display devices is recorded and displayed at low refresh rates (for example, 50 fields/sec for interlaced video material, and 24 frames/sec for film-originated material). One associated problem with such devices, including progressive digital display devices, is the presence of display artifacts. For example, one display artifact referred to as "area flicker" can occur due to the low display refresh rate. The area flicker becomes more visible as the size of the display increases, due to the high sensitivity to flicker in the human visual peripheral region. A simple solution for reducing the area flicker is to increase the display refresh rate by repeating the input fields or frames at a higher rate (for example, 100 fields/sec for interlaced video). This solves the area flicker problem for static scenes. However, the repetition introduces a new artifact in scenes with motion, known as "motion judder" or "motion smear," particularly in areas with high contrast, due to the human eye's tendency to track the trajectory of moving objects. For this reason, motion compensated frame interpolation is preferred, in which the pixels are computed in an interpolated frame or field at an intermediate point on a local motion trajectory, so that there is no discrepancy between an expected image motion due to eye tracking and a displayed image motion. The local image motion trajectory from one field or frame to the next is described by a motion vector.

[0003] Motion vectors can be computed at different levels of spatial resolution, such as at a pixel level, at an image patch level, or at an object level. "Image patch" refers to any portion of an image displayed in a frame. The image patch can be a single pixel, a plurality of pixels, and can have various shapes and sizes. Computing a motion vector for every pixel independently would theoretically result in an ideal data set, but is unfeasible due to the large number of computations required. Computing a motion vector for each image patch reduces the number of computations, but can result in artifacts due to motion vector discontinuities within an image patch. Computing motion vectors on an object basis can theoretically result in

high resolution and lower computational requirements, but object segmentation is a challenging problem.

[0004] Image noise and other problems can lead to errors in the computation and processing of motion vectors. Various techniques have been proposed in the search for accurate motion vector estimation. One of these techniques is the camera model, in which a mathematical model represents the movement of a camera which recorded the sequence of frames in a video signal. Camera models can provide mathematical representations of various camera movements including camera pans, zooms, and rotations. For instance, in a camera pan movement, the camera model provides a mathematical representation of the motion vectors associated with the camera moving in a horizontal and/or vertical direction at a constant velocity. The camera model is desirable because it can provide a global model of all of the motion vectors in an image patch or entire image frame. Thus, applying the camera model, every motion vector can be mathematically predicted at every location in the image frame.

[0005] One problem with conventional uses of camera models for motion vector estimation arises when there are two or more independently moving objects in a sequence of frames. In this situation, the independently moving objects introduce flaws into the camera model when attempting to fit the camera model to image data associated with the independently moving objects. For instance, in one sequence of frames, there are two moving objects: a car and a background image. The car, occupying 10 % of the frame, moves westward in the horizontal direction. The background image, occupying 90 % of the frame, moves eastward, opposite the car. The background image is the desired image for testing. Unless remedial measures are taken, the predicted camera model motion vectors associated with the background image will be erroneous due to the effect of the car movement. In particular, conventional techniques for computing the camera model would improperly fit a camera rotation to the background image rather than a proper camera pan. Resulting motion vectors predicted by this flawed camera model would also be erroneous at every point in the image.

[0006] Therefore, what is needed is a technique for fitting a camera model to a sequence of image frames wherein data associated with independently moving

objects other than a particular moving object or background to be tested is excluded to achieve a more accurate camera model.

SUMMARY

[0007] Aspects of the present invention involve defining and applying a camera model which provides a robust fit for predicting motion vectors associated with an independently moving object or background in a sequence of frames in which there may be other independently moving objects.

[0008] In general, in one aspect, the invention provides methods and apparatus, including computer program products, implementing and using techniques for camera pan vector estimation. A camera model is provided for representing motion vectors based on a plurality of parameters. A distribution of the motion vectors is generated. A cluster in the distribution of motion vectors is identified. A center of mass of the identified cluster is determined. The camera pan vector is provided based upon the determined center of mass.

[0009] Implementations can include one or more of the following features. The motion vectors can include a two-dimensional set of values, and generating the distribution of the motion vectors can include generating a two-dimensional histogram having bins. The motion vector values are provided in appropriate bins of the two-dimensional histogram. Also, before generating the distribution of the motion vectors, the motion vectors can be projected in a lower resolution.

[0010] Identifying the cluster in the distribution of motion vectors can include accumulating bins within a first neighborhood in the two-dimensional histogram, accumulating bins within a second neighborhood in the two-dimensional histogram, selecting one of the neighborhoods as having a higher accumulation of bins, and providing the selected neighborhood as the identified cluster in the distribution of motion vectors.

[0011] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0012]** FIG. 1 shows a flowchart of a recursive hierarchical process for determining a motion vector.
- [0013]** FIG. 2A shows an example of motion vectors associated with a camera pan.
- [0014]** FIG. 2B shows an example of motion vectors associated with a camera rotation.
- [0015]** FIG. 2C shows an example of motion vectors associated with a camera zoom in operation.
- [0016]** FIG. 3 shows a block diagram of an apparatus for calculating camera model parameters and generating motion vectors using a camera model.
- [0017]** FIG. 4 shows a flowchart of a method for estimating a camera pan vector, in accordance with one embodiment of the present invention.
- [0018]** FIG. 5 shows an example of motion vectors associated with independently moving objects.
- [0019]** FIG. 6A shows a distribution of motion vectors along two dimensions.
- [0020]** FIG. 6B shows a side view of a distribution of motion vector values.
- [0021]** FIG. 7 shows a distribution of motion vector values in a two-dimensional histogram having bins.
- [0022]** FIG. 8 shows a flowchart of a method for motion vector estimation using a camera model for a digital video sequence of frames.
- [0023]** FIG. 9 shows an apparatus for estimating a camera pan vector, in accordance with one embodiment of the present invention.
- [0024]** Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0025] Implementations of the present invention include a camera model used to provide a mathematical model of motion vectors for an image patch. The camera model represents a camera movement in a sequence of frames. For example, camera models can be used to provide for horizontal and/or vertical camera pans, zoom in or zoom out operations, and rotation operations. The camera model can also represent combinations of one or more of the above camera movements. The camera model provides a global description of motion vectors for a sequence of frames.

[0026] The camera model describes a global movement of the contents of the frame, as opposed to the local vectors at each image patch location that are computed independently and can therefore be used to aid in finding a more accurate, or true, motion vector. Typically, a motion vector resulting from camera movements at every location in a frame can be predicted quite easily with a simple model. For example, in the case of a camera lens panning across a distant landscape, all the motion vectors will be identical and equivalent to the velocity of the camera. Another scenario is when a camera lens zooms into an object on a flat surface, such as a picture on a wall. All the motion vectors then have a radial direction and increase from zero at the image center to a maximum value at the image periphery.

[0027] Embodiments of the present invention provide methods and apparatus for defining and applying a camera model which provides a robust fit for predicting motion vectors associated with an independently moving object in a sequence of frames in which there may be other independently moving objects. Disclosed herein are techniques for fitting a camera model to a sequence of image frames wherein data associated with independently moving objects other than a particular moving object or background to be tested is excluded to achieve a more accurate camera model. By excluding independently moving objects other than the particular moving object to be tested, a more accurate camera model is calculated, providing for more accurate motion vector prediction.

[0028] Embodiments of the present invention provide for the determination of parameters for a camera model, and the generating of motion vectors using the determined camera model parameters. A set of existing motion vectors stored in memory is provided, from which the camera model parameters are calculated. Applying a mathematical model, described below, a best fit method (such as linear least squares) is performed on the existing motion vectors to estimate the camera model parameters. The camera model parameters are then used to predict motion vectors for subsequent frames in a digital video sequence of frames. In a particular embodiment, this subsequent prediction of motion vectors using the camera model incorporates a recursive and hierarchical process such as that described in co-pending U.S. patent application entitled, "Temporal Motion Vector Filtering," having Application No. 10/833,198 filed April 26, 2004 which is incorporated by reference in its entirety for all purposes.

[0029] When calculating camera model parameters for a sequence of frames in which there are independently moving objects (some of which are moving in adverse directions to each other), embodiments of the present invention provide techniques for fitting the camera model to a dominant one of the independently moving objects and excluding the other independently moving objects. In most situations, the dominant object is the background. A set of input motion vectors having a first, or original resolution, are projected to a lower resolution to conserve memory resources. A histogram of all of the motion vector values is then generated to produce a distribution of the motion vectors. Motion vector values in bins of the histogram are accumulated to identify a dominant cluster of the motion vectors. Then a center of mass of this cluster is calculated while excluding other clusters of motion vectors in the two-dimensional distribution to achieve the desired robust fit. This center of mass provides one or more translation (or pan parameters) of the camera model. The camera model parameters are then used for later estimation of motion vectors using the camera model. In this way, motion vectors can be determined efficiently and accurately, such that little or no discrepancy exists between an expected image motion due to eye tracking and a displayed image motion in a digital video.

[0030] A mathematical representation of the camera model is as follows:

$$(1) \quad \begin{bmatrix} M_y \\ M_x \end{bmatrix} = \begin{bmatrix} S_y & r_1 \\ r_2 & S_x \end{bmatrix} \begin{bmatrix} y \\ x \end{bmatrix} + \begin{bmatrix} t_y \\ t_x \end{bmatrix}$$

[0031] In the equation above, the camera model provides an estimate of a plurality of motion vectors, M_y and M_x , by applying camera model parameters S_y , S_x , r_1 , r_2 , t_y , and t_x to a sequence of (x, y) locations in an image. S_y , S_x , r_1 , r_2 , t_y , and t_x are parameters of the camera model representing camera movements. In the above model, t_y and t_x represent a horizontal and vertical translation, respectively, during a camera pan. An illustration of motion vectors associated with a camera pan is provided in Fig. 2A. In the equation above, r_1 and r_2 are parameters representing camera rotation. Fig. 2B shows an illustration of motion vectors associated with a camera rotation. The S_y and S_x parameters are associated with a zoom out or zoom in operation which is illustrated in Fig. 2C.

[0032] In the described embodiment, the parameters S_x , S_y , r_1 , r_2 , t_y , and t_x are computed using a set of motion vectors determined using the recursive hierarchical process described in U.S. Patent Application No. 10/833,198. Accordingly, a motion vector is determined by applying a source correlation window to a first image frame and a target correlation window to a subsequent image frame, and placing the target correlation window such that a best match with the source correlation window is obtained, that is, the contents of the source correlation window and target correlation window are as similar as possible. In addition, the recursive hierarchical method uses multiple resolution levels of the image frames. A best motion vector is first determined for the lowest resolution level by projecting the previous best motion vector at the highest resolution level down to the lowest resolution level, and testing it and one or more updates. This best motion vector is then propagated up to a higher resolution level, where some adjustments are made and a new best motion vector is determined. This new best motion vector is propagated up to yet another higher resolution level, where more adjustments are made and another new best motion vector is determined. This process is repeated until the highest, original, resolution level has been reached and a best motion vector has been identified for the original resolution level.

[0033] FIG. 1 shows one implementation of the recursive hierarchical process (100) of U.S. Application No. 10/833,198. It is assumed that multiple resolution

levels of the image frames have already been generated. As can be seen in FIG. 1, the recursive hierarchical process (100) for determining a motion vector starts by projecting a motion vector from a previous image frame down to a lowest resolution level (step 102). A set of update vectors is generated and tested to find a best motion vector at this lowest resolution level (step 104). In one implementation this test is performed by comparing pixels in corresponding positions in a source correlation window centered on the origin of the motion vector and a target correlation window centered on the end point of each respective update vector. The comparison can, for example, be performed by subtracting a luma value for each pixel in the source window from the corresponding pixel in the respective target windows. In this case the best match would be defined by finding a minimum sum of absolute differences (SAD) for a source correlation window and a target correlation window pair, and the best motion vector would be the vector associated with this source correlation window and a target correlation window pair.

[0034] After the minimum SAD has been found, the best vector is selected (step 106). The process (100) then examines whether there are any higher resolution levels (step 108). If there are higher resolution levels, the process propagates the best vector up to the next higher resolution level (step 110) and repeats steps 104 through 108. If there are no higher resolution levels, the process proceeds to step 112, where the best vector is selected as the motion vector and is used for motion compensation, which completes the process for the current frame.

[0035] The best vector, as determined in step 112 of Fig. 1, can then be inserted into the mathematical model represented by equation (1) above as M_y and M_x . Those skilled in the art should appreciate that the best vector determined for any of the lower resolutions in Fig. 1, such as step 104, may also be used. In other embodiments, other vectors may be substituted for M_y and M_x in the camera model equation. It should be noted that although a linear least squares fit procedure is used in the described embodiment, any suitable data fitting technique can be used. Accordingly, a linear least squares fit procedure is then performed to solve for S_x , S_y , r_1 , r_2 , t_y , and t_x . That is, referring to the mathematical model above, since a known set of M_y and M_x values are provided for known corresponding x and y locations, a matrix inversion operation can be performed to determine the unknown parameters S_x , S_y , r_1 , r_2 , t_y , and t_x .

[0036] After the camera model parameters are determined, the camera model parameters are applied to generate camera model motion vectors M_y and M_x for positions y and x within an image patch or frame to provide vector candidates for one or more iterations of the recursive and hierarchical process described herein. For instance, these camera model predicted motion vectors can be used along with the updated vectors as test candidates for a SAD computation. Then, in step 106 of Fig. 1, the best vector is determined from all of the test candidates.

[0037] Fig. 3 shows a block diagram of an apparatus 300 for calculating camera model parameters and then using those camera model parameters to generate motion vectors in accordance with an embodiment of the invention. Fig. 4 shows a flow chart of a method 400 describing the operation of the apparatus of Fig. 3. In step 405, a test vector M_y' and M_x' , such as the best vector determined in step 106 of Fig. 1, is retrieved from memory 305 by a camera model parameter calculator 310. This camera model parameter calculator 310 is preferably implemented primarily in software, although various combinations of software and hardware can be used as will be understood by those skilled in the art. Memory 305 is any suitable processor readable storage medium such as SDRAM.

[0038] In step 410 of Fig. 4, the camera model parameter calculator 310 computes camera model parameters S_x , S_y , r_1 , r_2 , t_y , and t_x based on the retrieved motion vectors and known corresponding x,y locations, as described above. That is, a matrix inversion operation is performed to determine the camera model parameters S_x , S_y , r_1 , r_2 , t_y , and t_x . In step 415, camera model parameter calculator 310 then outputs these computed camera model parameters S_x , S_y , r_1 , r_2 , t_y , and t_x to memory 310 for storage in one or more registers. In step 420, a camera model motion vector generator 315 retrieves the camera model parameters S_x , S_y , r_1 , r_2 , t_y , and t_x from the registers in memory 310. Accordingly, when x,y locations for an image are provided to camera model motion vector generator 315 in step 425, the camera model motion vector generator 315 can generate motion vectors M_y and M_x in step 430, applying the camera model equation above. These motion vectors M_y and M_x are output by camera model motion vector generator 315 and can be stored in step 435. For the next iteration of the recursive hierarchical process, the stored motion vectors can be used as test candidates in addition to the recursive hierarchical predicted vectors.

[0039] In Fig. 5, a two-dimensional representation 500 of motion vectors along X and Y directions is illustrated. In Fig. 5, a group of motion vectors 505a-505e represent an independently moving object to which a camera model is desirably fitted. On the other hand, motion vectors 510 are motion vectors associated with a further independently moving image or object that is desirably excluded from the camera model. In this example, the camera model provides a pan scenario; thus, only the t_x and t_y values of the mathematical model above need be computed. Here, because the majority of motion vectors 505a-505e are the same, the camera pan model is desirably fitted to provide a model of these motion vectors 505a-505e. For the camera model to be effective, the motion vectors 510 are desirably excluded as outliers when calculating the t_x and t_y translation parameters for the camera pan model. Techniques described below provide for the exclusion of motion vectors 510 from the camera model parameter calculations, so that a robust and accurate camera model fit is achieved.

[0040] Fig. 6A shows a two-dimensional distribution 600 of M_y and M_x values for an exemplary set of motion vectors M_y and M_x . All of the values of the motion vectors M_y and M_x are within the bounds of the two-dimensional representation. That is, all of the M_x values range within $-M_x \text{ max}$ and $+M_x \text{ max}$; similarly, all of the M_y values range within $-M_y \text{ max}$ and $+M_y \text{ max}$. In one example, the M_x values range between -32 and $+32$, and the M_y values range between -24 and $+24$. Those skilled in the art should appreciate that these ranges of M_y and M_x values are only exemplary, and various other ranges may be used in accordance with embodiments of the present invention. In Fig. 6, the majority of the motion vectors form a cluster 605, indicating that all of the motion vectors have substantially similar values representing a camera pan. Thus, determining the center of mass of the cluster 605 provides the t_x and t_y parameters for the camera model. The vectors outside of cluster 605 are preferably excluded as outliers during the computation of t_x and t_y .

[0041] Fig. 6B shows a side view of a distribution of M_x values, such as the motion vector values of Fig. 6A. The distribution of M_x values yields a first hill 615 and a second hill 620. The hills 615 and 620 indicate the number of vectors having particular M_x values. Thus, first hill 615 is formed because of a large concentration of vectors having the value M_{x1} . The relatively smaller second hill 620 is formed by another concentration of vectors having the value M_{x2} . The hill 615 having the larger

count of vectors for a particular M_x value, referred to herein as the dominant hill, represents a neighborhood for locating the true t_x and t_y parameters for the camera model. Methods and apparatus according to embodiments of the present invention provide for identifying the dominant hill 615, excluding smaller hills such as hill 620, then determining a center of mass of the dominant hill 615 as the camera model parameters t_x and t_y .

[0042] Fig. 7 shows a top view of an example of a two-dimensional histogram with bins 705a-705c for accumulating discrete M_y and M_x values. Because the M_y and M_x values are discrete, the histogram is represented as a grid, with each block or bin in the grid representing a particular M_y and M_x value. In the example in which the M_x values range between -32 and $+32$, and the M_y values range between -24 and $+24$, a two-dimensional distribution of the motion vector values would have a grid size of 64×48 total distributions or bins in which the particular motion vector values can be located. The aggregation of the bins provides a 64×48 two-dimensional histogram. The bins 705 have bin counts indicating the number of motion vector M_y and M_x values in the respective bins. For instance, bin 705b has a count of 7 motion vectors, and bin 705c has a count of 100 motion vectors. Each bin count is incremented for every motion vector value deposited in that bin. In this example, a neighborhood of bins including bin 705c represents the dominant hill of the histogram. It is desirable to identify this dominant hill for determining an accurate camera pan vector estimate using techniques described below.

[0043] Figs. 8 and 9 show a method 800 and apparatus 900, respectively, for estimating a camera pan vector, performed in accordance with one embodiment of the present invention. The apparatus 900 can be implemented with various combinations of hardware and software, as will be understood by those skilled in the art. The method of Fig. 8 is described with reference to Fig. 9. In step 805, a camera model is provided. This camera model is defined according to the method described with reference to Figs. 3 and 4, as explained above. The camera model is used to generate a plurality of camera model motion vectors, for instance, a two-dimensional set of motion vectors M_y and M_x , as explained above, for an image patch in a plurality of frames of the digital video sequence. The camera model includes a plurality of parameters, such as t_x and t_y , as described above, which are

applied to y and x coordinates to generate the motion vectors, M_y and M_x . The camera model motion vectors are provided in a first resolution, such as 1:1.

[0044] In step 810 of Fig. 8, the camera model motion vectors M_y and M_x are preferably projected in a lower resolution. This step of projecting the motion vectors in the lower resolution conserves memory and reduces computational demands. Various resolutions may be used. In one example, where the original or first resolution is 1:1, this lower resolution is 1:2, 1:4, or 1:8. Accordingly, the number of bins in the two-dimensional histogram of Fig. 9 is greatly reduced. For instance, when the original M_x value range of -32 to +32 and value range of -24 to +24 is reduced to 1:2, the corresponding bin count is reduced from 64x48 to 32x24. That is, scaling the motion vector values by a factor of 2 causes the bin count to be reduced by a factor of 4. In this example, the bin counts for 4 neighboring bins are effectively combined into a single bin. Various other reduced resolutions may be used, as will be understood by the skilled artisan.

[0045] In step 815 of Fig. 8, a motion vector generator 910, shown in Fig. 9, generates a distribution of the motion vectors. When the motion vectors are two-dimensional, the particular values of the set of two-dimensional motion vectors are distributed in a two-dimensional histogram with bins, as explained above with reference to Fig. 7.

[0046] In step 820 of Fig. 8, a cluster identifier 920, shown in Fig. 9, finds the dominant hill in the two-dimensional histogram of Fig. 7. That is, cluster identifier 920 identifies a cluster in the distribution of motion vector values as having the highest concentration of motion vectors in the distribution. In one embodiment, this identification step includes accumulating bins within a plurality of neighborhoods in the two-dimensional histogram as shown in Fig. 7. These neighborhoods are defined by a $p \times q$ window 710 or other window of predetermined size, e.g. 3 x 3, 4 x 4, 12 x 12, 100 x 100, moved over various positions in the histogram such that a certain number of bins are included within the window at each position. All of the bin counts within the window at a certain position are added to provide a total bin count associated with that window or neighborhood. Then, the neighborhood having the higher or highest accumulation of bins is selected as the identified cluster. This identified cluster provides the dominant hill of the histogram, such as hill 615 in Fig. 6B.

[0047] In step 825 of Fig. 8, a determining module 930 shown in Fig. 9 determines a center of mass for the identified cluster from step 1020. This center of mass calculation is preferably performed within another window of predetermined size. In one embodiment, this window has a size $m \times n$ which is larger than the $p \times q$ window containing the neighborhoods used for the bin accumulations in step 820. In step 830, the center of mass determined in step 825 is provided as one or more of the camera model parameters for subsequent motion vector calculation, referring to the mathematical representation of the camera model above. For instance, in the camera pan example, the determined center of mass is provided as the translation parameters, t_x and t_y . In embodiments in which the motion vectors were projected from an original resolution to a lower resolution in step 810, the determined center of mass is preferably projected back to the original resolution. In some embodiments, the calculations for this projection are done using floating point to provide a more accurate projection. The camera model parameters are then stored in step 835, and later retrieved by camera pan calculator 940 of Fig. 9 for the generation of camera pan model motion vectors M_y and M_x , using equation (1) above. In one embodiment, these camera pan model motion vectors M_y and M_x are then used to determine test candidates for the recursive hierarchical process described in U.S. Application No. 10/833,198.

[0048] Using the camera model predicted vector as an additional candidate vector in the recursive hierarchical process minimizes the possibility that the recursive portion of the recursive hierarchical search converges into a false local minimum instead of the true minimum. The camera predicted vector candidate can direct the process towards a true minimum. In some embodiments, as described herein, motion vectors predicted or generated by the camera model are also used as test candidates for SAD computation.

[0049] Embodiments of the invention, including the apparatus of Figs. 3 and 9, can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Apparatus embodiments of the invention can be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and method steps of the invention can be performed by a programmable processor executing a program of instructions to perform functions of the invention by operating on input

data and generating output. Embodiments of the invention can be implemented advantageously in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. Each computer program can be implemented in a high-level procedural or object-oriented programming language, or in assembly or machine language if desired; and in any case, the language can be a compiled or interpreted language. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory and/or a random access memory. Generally, a computer will include one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

[0050] A number of implementations of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

CLAIMS

1. A method for estimating a camera pan vector comprising:

generating a distribution of motion vectors;

5 identifying a cluster in the distribution of motion vectors;

determining a center of mass of the identified cluster; and

providing the camera pan vector based upon the determined center of mass.

10 2. The method of claim 1, the camera pan vector including parameters of a camera model representing the motion vectors in an image patch.

3. The method of claim 2, wherein providing the camera pan vector based upon the determined center of mass comprises:

15 providing the determined center of mass as one or more of the parameters of the camera model.

4. The method of claim 1, further comprising:

storing the provided camera pan vector in a processor readable memory.

20 5. The method of claim 1, the motion vectors including values, wherein generating the distribution of the motion vectors comprises:

generating a two-dimensional histogram having bins, and

providing the motion vector values in appropriate bins of the two-dimensional histogram.

25

6. The method of claim 1 further comprising:

providing the motion vectors in a first resolution, and

before generating the distribution of the motion vectors,
projecting the motion vectors in a second resolution, the second resolution of lower resolution than the first resolution.

5 7. The method of claim 1, wherein the camera pan vector is generated for an image patch, the image patch including a single pixel.

8. The method of claim 1, wherein the camera pan vector is generated for an image patch, the image patch including a plurality of pixels.

10

9. The method of claim 5, wherein identifying the cluster in the distribution of motion vectors comprises:

accumulating bins within a first neighborhood in the two-dimensional histogram,

15 accumulating bins within a second neighborhood in the two-dimensional histogram,

selecting one of the neighborhoods as having a higher accumulation of bins, and

20 providing the selected neighborhood as the identified cluster in the distribution of motion vectors.

10. The method of claim 9, wherein each neighborhood is contained within a window of a predetermined size.

25 11. The method of claim 9, wherein identifying the cluster in the distribution of motion vectors further comprises:

before selecting the one neighborhood,

accumulating bins within further neighborhoods in the two-dimensional histogram.

12. The method of claim 1, wherein the center of mass of the identified cluster is determined within a window of predetermined size.

5 13. The method of claim 10, wherein the center of mass of the identified cluster is determined within a window having a larger size than the window or windows containing the neighborhoods.

14. An apparatus for estimating a camera pan vector comprising:

10 a motion vector generator configured to generate a distribution of motion vectors;

a cluster identifier configured to identify a cluster in the distribution of motion vectors;

15 a determining module configured to determine a center of mass of the identified cluster; and

a camera pan calculator configured to provide the camera pan vector based upon the determined center of mass.

20 15. The apparatus of claim 14, the camera pan vector including parameters of a camera model representing the motion vectors in an image patch.

16. The apparatus of claim 15, the camera pan calculator configured to:

provide the determined center of mass as one or more of the parameters of the camera model.

25

17. The apparatus of claim 14, the motion vectors including values, the motion vector generator configured to:

generate a two-dimensional histogram having bins, and

provide the motion vector values in appropriate bins of the two-dimensional histogram.

18. The apparatus of claim 17, the cluster identifier configured to:

5 accumulate bins within a first neighborhood in the two-dimensional histogram,
 accumulate bins within a second neighborhood in the two-dimensional histogram,

 select one of the neighborhoods as having a higher accumulation of bins, and

 provide the selected neighborhood as the identified cluster in the distribution
10 of motion vectors.

19. A computer program product, stored on a processor readable medium, comprising instructions operable to cause a computer to perform a method for estimating a camera pan vector, the method comprising:

15 generating a distribution of motion vectors;

 identifying a cluster in the distribution of motion vectors;

 determining a center of mass of the identified cluster; and

 providing the camera pan vector based upon the determined center of mass.

20 20. The computer program product of claim 19, the motion vectors including values, wherein generating the distribution of the motion vectors comprises:

 generating a two-dimensional histogram having bins, and

 providing the motion vector values in appropriate bins of the two-dimensional histogram.

25

21. The computer program product of claim 20, the step of identifying the cluster in the distribution of motion vectors including:

accumulating bins within a first neighborhood in the two-dimensional histogram,

accumulating bins within a second neighborhood in the two-dimensional histogram,

5 selecting one of the neighborhoods as having a higher accumulation of bins,
and

providing the selected neighborhood as the identified cluster in the distribution of motion vectors.

10

**ROBUST CAMERA PAN VECTOR ESTIMATION USING ITERATIVE CENTER OF
MASS**

5

Methods and apparatus, including computer program products, implementing and using techniques for camera pan vector estimation, are disclosed. A camera model is provided for representing motion vectors based on a plurality of parameters. A distribution of the motion vectors is generated. A cluster in the distribution of

10 motion vectors is identified. A center of mass of the identified cluster is determined.

○ The camera pan vector is provided based upon the determined center of mass.

Figure 8