UC San Diego

UC San Diego Previously Published Works

Title

Motion analysis for event detection and tracking with a mobile omnidirectional camera

Permalink

<https://escholarship.org/uc/item/1qw82651>

Journal

Multimedia Systems, 10(2)

ISSN

0942-4962

Authors

Gandhi, Tarak
Trivedi, Mohan M

Publication Date

2004-08-01

Peer reviewed

Motion Analysis for Event Detection and Tracking with a Mobile Omni-Directional Camera

Tarak Gandhi
Computer Vision and Robotics Research
Laboratory
University of California – San Diego
tgandhi@ucsd.edu

Mohan M. Trivedi
Computer Vision and Robotics Research
Laboratory
University of California – San Diego
mtrivedi@ucsd.edu

ABSTRACT

A mobile platform mounted with Omni-Directional Vision Sensor (ODVS) can be used to monitor large areas and detect interesting events such as independently moving persons and vehicles. To avoid false alarms due to extraneous features, the image motion induced by the moving platform should be compensated. This paper describes a formulation and application of parametric ego-motion compensation for an ODVS. Omni images give 360 degrees view of surroundings but undergo considerable image distortion. To account for these distortions, the parametric planar motion model is integrated with the transformations into omni image space. Prior knowledge of approximate camera calibration and camera speed are integrated with the estimation process using Bayesian approach. Iterative, coarse to fine, gradient based estimation is used to correct the motion parameters for vibrations and other inaccuracies in prior knowledge. Experiments with camera mounted on various types of mobile platforms demonstrate successful detection of moving persons and vehicles.

Keywords

Motion detection, Optical flow, Panoramic vision, Dynamic vision, Mobile robots, Intruder detection, Surveillance

1. INTRODUCTION AND MOTIVATION

Computer vision researchers have for long recognized the importance of visual surveillance related applications while pursuing some of the outstanding research issues in dynamic scene analysis, motion detection, feature extraction, pattern and activity analysis and biometric systems. Recent world events are demanding practical and robust deployment of video based solutions for a wide range of applications [11, 16, 30, 33]. Such wider acceptance of the need of the technology does not mean that these systems are indeed ready for deployment. There are many important and difficult research problems that are still outstanding. In this paper we

present a research study focused on one of such challenging research problem, that of developing an autonomous system which can serve as a “Mobile Sentry” to perform the tasks that a person posted on a guard duty around a perimeter of a base gets assigned. The mobile sentry with video cameras should be able to detect “interesting” events, record and report the nature and location of the event in real-time for further processing by a human operator. This is indeed an ambitious goal and it requires resolution of several important problems from computer vision and intelligent robotics area. In this paper we are focused on the problem of how to detect and compensate for the ego motion of the mobile platform. One of the novel feature of our research is the omnidirectional video streams we use as the input to the vision system.

When the camera is stationary, background subtraction is often used to extract moving objects [20, 39]. However, when the camera is moving, the background also undergoes ego-motion, which should be compensated. To distinguish objects of interest from extraneous features on the ground, the ground is usually approximated by a planar surface, whose ego-motion is modelled using a projective transform [12, 29] or its linearized version. Using this model, the ego-motion of the ground can be compensated in order to separate the objects with independent motion or height. This approach has been widely used for object detection from moving platforms [6, 13, 29].

Omni-Directional Vision Sensors (ODVS) or omni-cameras which give a 360 degree field of view of the surroundings are very useful for applications such as surveillance [5, 8], robot navigation [42], localization [26], and wide baseline stereo [38, 43]. The book by Benosman [3] gives a comprehensive review of the theory and applications of omni cameras. Motion estimation from moving ODVS cameras has recently been a topic of great interest. Rectilinear cameras usually have a smaller field of view, due to which the focus of expansion often lies outside the image, causing motion estimation to be sensitive to the camera orientation. Also, the motion field produced by translation along horizontal direction is similar to that due to rotation about vertical axis. As noted by Gluckman and Nayar [18], ODVS cameras avoid both these problems due to their wide field of view. They project the image motion on a spherical surface using Jacobians of transformations to determine ego-motion of a moving platform terms of translation and rotation of the camera.

Vassalo et al. [41] use the spherical projection to determine ego-motion of a moving platform terms of translation and rotation of the camera. Experiments are performed using robotic platforms in indoor environment and the ego-motion estimates are compared with those from odometry. Shakerinia et al. [35] use the concept of back-projection flow, where the image motion is projected to a virtual curved surface in place of spherical surface and make the Jacobians of transformation simpler. Using this concept, they have adapted ego-motion algorithms for use with ODVS sensors. Results using simulated image sequences show the basic feasibility of the approach.

In our own research, the emphasis is on robustness, efficiency, and applicability in outdoor environments encountered in surveillance and physical or base security. The main contribution of this paper is to perform detection of events such as independently moving persons and automobiles from ODVS image sequences, and apply it to video sequences obtained from a moving platform for surveillance applications.

2. EGO-MOTION COMPENSATION FRAMEWORK FOR ODVS VIDEO

Parametric motion estimation based on image gradients, also known as the “direct method” has been used for rectilinear cameras for planar motion estimation, obstacle detection and motion segmentation [24, 28]. The advantage of the direct methods is that they can use motion information not only from corner-like features, but also edges, which are usually more numerous in an image. On the other hand, direct methods are more challenging for implementation, especially in outlier removal, and it is more difficult to track over frames. The comparison between the corner-based and direct gradient based methods is given in Table 1

Here, the concept of direct method is extended for use with ODVS. This approach was also used for detecting surrounding vehicles from a moving car in [15, 22]. An ODVS gives full 360 degree view of the surroundings, which reduces the motion ambiguities often present in the rectilinear cameras. However, the images undergo considerable distortion which should be accounted for during motion estimation.

The block diagram of the event detection system is shown in Figure 1. The initial estimates of the the ground plane motion parameters are obtained using the approximate knowledge about the camera calibration and speed. Using these parameters, one of the frames is warped towards another to compensate the motion of the ground plane. However, the motion of features having independent motion or height above the ground plane is not fully compensated. To detect these features, the normalized image difference between the two images is computed using temporal and spatial gradients. This suppresses the features on ground plane and enhances the interesting objects. Morphological and other post-processing is performed to further suppress the ground features due to any residual motion and get the positions of the objects. The detected objects are then tracked over frames to form events.

However, the calibration and speed of the camera may not be known accurately. Furthermore, the camera could vibrate during the motion. Due to this, the motion of the ground

Figure 1: Event detection and recording system based on ego-motion compensation from a moving platform: Block Diagram

plane may not be fully compensated, leading to misses and false alarms. In order to improve the performance, motion parameters are iteratively corrected using the spatial and temporal gradients of the motion compensated images using optical flow constraint in coarse to fine framework. The motion information contained in these gradients is optimally combined with the prior knowledge of motion parameters using a Bayesian framework similar to [29]. Robust estimation is used to separate the ground plane features from other features. The following sections deal with the individual blocks described above, along with the appropriate formulation for ODVS.

3. ODVS MOTION TRANSFORMATIONS

To compensate the motion of the omni-directional camera, the transformation due to ODVS should be combined with that due to motion. These transforms are discussed below:

3.1 Flat plane transformation

The ODVS used in this work consists of a hyperbolic mirror and a camera placed on its axis. It belongs to a class of cameras known as central panoramic catadioptric cameras [3]. These cameras have a single viewpoint that permits the image to be suitably transformed to obtain perspective views. Figure 2 (a) shows a photograph of an ODVS mirror.

Corner based methods	Gradient based (Direct) methods
Determines motion of individual features	Fits motion model to entire or part of scene
Only corner-like features used	Edge features used (more numerous)
Easier to track over frames	More difficult to track over frames
Easier to identify outliers	Outlier removal more difficult
Easier to implement	More difficult to implement

Table 1: Comparison between corner based and gradient based motion estimation

Figure 2: (a) Omni-Directional Vision Sensor (ODVS) (b) A typical image from the ODVS. (c) Transformation to a perspective plan view.

An image from a camera mounted with an ODVS mirror is shown in Figure 2 (b). It is seen that the camera covers a 360 degrees field of view around its center. However, the image it produces is distorted with straight lines transformed into curves. A flat plane transformation is applied to the image to produce a perspective view looking downwards as shown in Figure 2 (c), where the distortion is considerably reduced. Details of this transformation are discussed below.

The geometry of a hyperbolic ODVS is shown in Figure 3. According to the mirror geometry, a light ray from the object towards the viewpoint at the first focus O is reflected so that it passes through the second focus, where a conventional rectilinear camera is placed. The equation of the hyperboloid is given by:

$$\frac{(Z-c)^2}{a^2} - \frac{X^2 + Y^2}{b^2} = 1$$

where $c = \sqrt{a^2 + b^2}$.

Let $P = (X, Y, Z)^T$ denote the homogenous coordinates of the perspective transform of any 3-D point λP on ray OP , where λ is the scale factor depending on the distance of the 3-D point from the origin. It can be shown [1, 23, 35] that the reflection in mirror gives the point $-p = (-x, -y)^T$ on the image plane of the camera using the flat-plane transform F :

$$F(P) = p = \begin{pmatrix} x \\ y \end{pmatrix} = \frac{q_1}{q_2 Z + q_3 \|P\|} \begin{pmatrix} X \\ Y \end{pmatrix} \quad (1)$$

Figure 3: Omni-directional camera geometry

where

$$q_1 = c^2 - a^2, q_2 = c^2 + a^2, q_3 = 2ac$$

$$\|P\| = \sqrt{X^2 + Y^2 + Z^2}$$

Note that the expression for image coordinates p is independent of the scale factor λ . The pixel coordinates $w = (u, v)^T$ are then obtained by using the calibration matrix K of the conventional camera composed of the focal lengths f_u, f_v , optical center coordinates $(u_0, v_0)^T$, and camera skew s .

$$\begin{pmatrix} w \\ 1 \end{pmatrix} = K \begin{pmatrix} p \\ 1 \end{pmatrix}$$

or

$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \begin{pmatrix} f_u & s & u_0 \\ 0 & f_v & v_0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \quad (2)$$

This transform can be used to warp an omni image to a plan perspective view. To convert a perspective view back to omni view, the inverse flat-plane transform p can be used:

$$\begin{pmatrix} p \\ 1 \end{pmatrix} = \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = K^{-1} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \quad (3)$$

$$F^{-1}(p) = P = \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} q_1 x \\ q_1 y \\ q_2 - q_3 \sqrt{x^2 + y^2 + 1} \end{pmatrix} \quad (4)$$

It should be noted that the transformation of omni to perspective view involves very different magnifications in different parts of the image. Due to this, the quality of the image deteriorates if the entire image is transformed at a time. Hence, it is desirable to perform motion estimation

directly in the ODVS domain, but use the above transformations to map the locations to the perspective domain as required.

3.2 Planar motion transformation

To detect objects with motion or height, the motion of the ground is modeled using planar motion model [12, 27]. Let P_A and P_B denote the perspective transforms of a point on the ground plane in the homogenous coordinate systems corresponding to two positions A and B of the moving camera. These are related by:

$$\lambda_B P_B = \lambda_A R P_A + D = \lambda_A [R P_A + D/\lambda_A] \quad (5)$$

where R and D denote the rotation and translation between the camera positions, and λ_A, λ_B depend on the distance of the actual 3-D point. Let the ground plane satisfy the following equation at the camera position A :

$$\lambda_A K^T P_A = 1$$

or

$$1/\lambda_A = K^T P_A$$

Substituting the value of $1/\lambda_A$ in equation (5), it is seen that P_A and P_B are related by a projective transform:

$$\lambda_B P_B = \lambda_A [R + D K^T] P_A = \lambda_A H P_A \quad (6)$$

or $P_B \equiv H P_A$ within a scale factor. This relation has been widely used to estimate planar motion for perspective cameras.

For performing motion compensation using omni-directional cameras, the above projective transform should be combined with the flat-plane transform as well as camera calibration matrix to warp every point in one image towards another. The complete transform for warping is shown in Figure 4.

4. PARAMETRIC MOTION ESTIMATION FOR ODVS

This section describes the main contribution of the paper. Direct methods based on image gradients have been applied for estimating the motion parameters for rectilinear cameras [4, 24]. Here, the direct method is generalized for ODVS cameras. Information from image gradients is combined with the a-priori known information about the camera motion and calibration in Bayesian framework to obtain optimal estimates of motion parameters for ego-motion compensation.

4.1 Use of optical flow constraint

Under favorable conditions, the spatial gradients (g_u, g_v) , the temporal gradient g_t , and the residual image motion $(\Delta u, \Delta v)^T$ after current motion compensation satisfy the optical flow constraint [21].

$$g_u \Delta u + g_v \Delta v + g_t = 0 \quad (7)$$

However, there is only one equation between two unknowns for each point. Due to this, only the normal flow - i.e., flow in the direction of the gradient - can be determined using a single point. This is known as the aperture problem, illustrated in Figure 5 (a). To solve this problem, Lucas

Figure 4: Transforming a pixel from omni image A to omni image B using (1) Inverse calibration matrix K^{-1} , (2) Inverse flat plane transform F^{-1} (3) Projective transform H for planar motion from A to B (4) Flat plane transform F (5) Calibration matrix K

and Kanade [31] assumed that the image motion is approximately constant in a small window around every point. Using this constraint, more equations are obtained using the neighboring points, and the full optical flow can be estimated using least squares. Such an estimate is reliable near corner-like points where window has gradients in different directions. This is as seen in Figure 5 (b). This method has been used by Kanade, Lucas and Tomasi [36] to find and track corner-like features over an image. However, in case of ODVS the assumption of uniform optical flow needs to be modified due to the non-linear ODVS transform. Daniilidis [9] has generalized the optical flow estimation to ODVS camera.

However, this approach would use the motion information

Figure 5: Aperture problem. (a) In the case of an edge, only the component of motion normal to the edge can be determined. (b) In the case of a corner, the aperture problem is avoided, and the motion can be uniquely determined.

only at corner-like features. However, the edge features also have motion information. To use this information, the image gradients can be used directly to estimate the model parameters. This approach is known as direct method of motion estimation in literature and has been extensively used in obstacle detection using rectilinear cameras [4, 24]. Usually, a linearized version of projective transform is used:

$$\begin{aligned}\Delta u &= a_1 u + a_2 v + a_3 + a_7 u^2 + a_8 uv \\ \Delta v &= a_4 u + a_5 v + a_6 + a_7 uv + a_8 v^2\end{aligned}$$

The expressions of image motion are substituted into the optical flow constraint in equation (7) to give:

$$g_u(a_1 u + a_2 v + a_3 + \dots) + g_v(a_4 u + a_5 v + a_6 + \dots) + g_t = 0$$

This gives one equation for every point in 8 parameters, that can be solved using linear least squares. Since the quadratic parameters are more sensitive to noise, a 6-parameter affine model is also used.

4.2 Generalization for ODVS

To apply the motion estimation to ODVS camera, the non-linear flat-plane transform is used to go from omni to perspective domain and back. Since non-linearity has to be dealt with anyway, the projective transform H is used instead of a linear model so that large motions can be handled better. The motion parameters in the projective transform are parameterized as:

$$\mathbf{h} = (h_1 \ h_2 \ h_3 \ h_4 \ h_5 \ h_6 \ h_7 \ h_8)^T$$

with

$$\frac{H}{H_{33}} = \begin{pmatrix} h_1 & h_2 & h_3 \\ h_4 & h_5 & h_6 \\ h_7 & h_8 & 1 \end{pmatrix}$$

The optical flow constraint equation is satisfied only for small image displacements up to 1 or 2 pixels. To estimate larger motions, a coarse to fine pyramidal framework [25, 37] is used. In this framework, a multi-resolution Gaussian pyramid is constructed for adjacent images in the sequence. The motion parameters are first computed at the coarsest level, and the image points at the next finer level are warped using the computed motion parameters. The residual motion is computed at the finer level, and the process is repeated until the finest level. Even within each level, multiple iterations of warping and estimation can be performed.

Let $\bar{\mathbf{h}}$ be the actual value of the motion parameter vector, and $\hat{\mathbf{h}}$ be the current estimate. Using the current estimate, the second image B is warped towards the first image A to get the warped image B' . Then, the transformation between A and B' can be expressed approximately in terms of $\Delta \mathbf{h} = \bar{\mathbf{h}} - \hat{\mathbf{h}}$. Let $w_A = (u_A, v_A)^T$ be the projection of a point on the planar surface in image A . Then, the projection w_B of the same point in warped image B' is a function of w_A as well as $\Delta \mathbf{h}$, given using a composition of operations shown in Figure 4. The optical flow constraint between images A and B' is then given by:

$$(g_u \ g_v) [w_{B'} - w_A] = -g_t + \eta$$

where η accounts for the random noise in the temporal image gradient. For N points on the planar surface, the constraints

can be expressed in a matrix form:

$$\Delta \mathbf{z} = \mathbf{c}(\Delta \mathbf{h}) + \mathbf{v}$$

where every row i of the equation represents the constraint for a single image point with

$$\begin{aligned}\mathbf{c}_i(\Delta \mathbf{h}) &= (g_u \ g_v)_i [w_{B'}(w_A; \Delta \mathbf{h}) - w_A] \\ \Delta \mathbf{z}_i &= -(g_t)_i, \mathbf{v}_i = \eta_i\end{aligned}\quad (8)$$

Due to the flat plane and the projective transforms, the function $\mathbf{c}(\cdot)$ is non-linear. Hence, state estimate $\hat{\mathbf{h}}$ and its covariance $\mathbf{P}$ are iteratively updated using the measurement update equations of the iterated extended Kalman filter [2], with $\mathbf{C}$ denoting the Jacobian matrix of $\mathbf{c}(\cdot)$.

$$\mathbf{P} \leftarrow [\gamma \mathbf{C}^T \mathbf{R}^{-1} \mathbf{C} + \mathbf{P}^{-1}]^{-1} \quad (9)$$

$$\hat{\mathbf{h}} \leftarrow \hat{\mathbf{h}} + \Delta \hat{\mathbf{h}} = \hat{\mathbf{h}} + \mathbf{P} [\gamma \mathbf{C}^T \mathbf{R}^{-1} \Delta \mathbf{z} - \mathbf{P}^{-1}(\hat{\mathbf{h}} - \mathbf{h}_-)] \quad (10)$$

where $\mathbf{R}$ is the covariance of the temporal gradient measurements, $\mathbf{h}_-$ is the prior value of state obtained from camera calibration and velocity, and $\mathbf{P}_-$ is the prior covariance. The matrix $\mathbf{R}$ is taken as a diagonal matrix to simplify calculations. However, this would mean assuming that the pixel gradients are independent, which may not really be the case since gradients are computed from multiple pixels. Hence, the factor $\gamma \leq 1$ is used to accommodate the interdependence of the gradient measurements.

To compute the Jacobian $\mathbf{C}$, each row $\mathbf{C}_i$ is expressed using chain rule:

$$\mathbf{C}_i = \left(\frac{\partial \mathbf{c}_i}{\partial \mathbf{h}} \right) = \left(\frac{\partial \mathbf{c}}{\partial w_B} \frac{\partial w_B}{\partial p_B} \frac{\partial p_B}{\partial P_B} \frac{\partial P_B}{\partial \mathbf{h}} \right)_i \quad (11)$$

where $P_B = (X_B, Y_B, Z_B)^T$, $p_B = (x_B, y_B)^T$ and $w_B = (u_B, v_B)^T$ are the coordinates of point i in the mirror, image, and pixel coordinate systems for camera position B .

Differentiating equation (8) w.r.t. w_B gives:

$$\left(\frac{\partial \mathbf{c}}{\partial w_B} \right)_i = (g_u \ g_v)_i$$

The calibration equation (2) can be differentiated to obtain:

$$\left(\frac{\partial w_B}{\partial p_B} \right)_i = \begin{pmatrix} f_u & s \\ 0 & f_v \end{pmatrix}$$

The Jacobian of the flat plane transform is obtained by differentiating equation (1) at $P = P_B$ as:

$$\begin{aligned}\left(\frac{\partial p_B}{\partial P_B} \right)_i &= \begin{pmatrix} \frac{\partial x_B}{\partial X_B} & \frac{\partial x_B}{\partial Y_B} & \frac{\partial x_B}{\partial Z_B} \\ \frac{\partial y_B}{\partial X_B} & \frac{\partial y_B}{\partial Y_B} & \frac{\partial y_B}{\partial Z_B} \end{pmatrix}_i \\ &= \frac{1}{(q_2 Z_B + q_3 \|P_B\|)_i \|P_B\|_i} \\ &\quad \cdot \begin{pmatrix} q_3 x_B X_B - q_1 \|P_B\| & q_3 x_B Y_B & q_3 x_B Z_B \\ q_3 y_B X_B & q_3 y_B Y_B - q_1 \|P_B\| & q_3 y_B Z_B \end{pmatrix}_i\end{aligned}$$

Since the ODVS transforms giving p_A and p_B do not change if the homogenous coordinates P_A and P_B are changed by scale factor, we can scale the right hand side of equation (6) to give:

$$P_B = \frac{1}{H_{33}} H P_A = \begin{pmatrix} h_1 X_B + h_2 Y_B + h_3 Z_B \\ h_4 X_B + h_5 Y_B + h_6 Z_B \\ h_7 X_B + h_8 Y_B + Z_B \end{pmatrix}$$

Taking the Jacobian w.r.t. $\mathbf{h} = (h_1 \dots h_8)$ gives:

$$\left(\frac{\partial P_B}{\partial \mathbf{h}} \right)_i = \begin{pmatrix} X_B & Y_B & Z_B & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & X_B & Y_B & Z_B & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & X_B & Y_B \end{pmatrix}_i$$

4.3 Outlier removal

The estimate given above is optimal only when all points really belong to the planar surface, and the underlying noise distributions are Gaussian. However, the estimation is highly sensitive to the presence of outliers, i.e. points not satisfying the ground motion model. These features should be separated using a robust method. To reduce the number of outliers, the region of interest of road is determined using calibration information, and the processing is done only in that region to avoid extraneous features. To detect outliers an approach similar to the data snooping approach discussed in [10] has been adapted for Bayesian estimation. In this approach, the error residual of each feature is compared with the expected residual covariance at every iteration, and the features are reclassified as inliers or outliers.

If a point $\mathbf{z}_i$ is not included in the estimation of $\hat{\mathbf{h}}$ – i.e. is currently classified as an outlier – then the covariance of its residual is:

$$\text{Cov} [\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] \simeq \mathbf{R} + \mathbf{C}_i \mathbf{P} \mathbf{C}_i^T$$

However, if $\mathbf{z}_{in}$ is included in the estimation of $\hat{\mathbf{h}}$ – i.e. is currently classified as an inlier – then it can be shown that the covariance of its residual is given by:

$$\text{Cov} [\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] \simeq \mathbf{R} - \mathbf{C}_i \mathbf{P} \mathbf{C}_i^T < \mathbf{R}$$

Hence, to classify in the next iteration, the Mahalanobis norm of the residual is compared with a threshold τ . For point currently classified as inlier the following condition is used:

$$[\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] [\mathbf{R} + \mathbf{C}_i \mathbf{P} \mathbf{C}_i^T]^{-1} [\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] < \tau \quad (12)$$

For point currently classified as inlier the covariance $\mathbf{R}$ is used in practice instead of $\mathbf{R} - \mathbf{C}_i \mathbf{P} \mathbf{C}_i^T$ in order to avoid non-positive definite covariance because of approximations due to non-linearities. This would somewhat increase the probability of classifying as an outlier instead of inlier, which is to be on the safer side.

$$[\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] \mathbf{R}^{-1} [\Delta \mathbf{z}_i - \mathbf{C}_i \Delta \hat{\mathbf{h}}] < \tau \quad (13)$$

Note that this method is effective only when there is some prior knowledge about the motion parameters, otherwise the prior covariance $\mathbf{P}_-$ would become infinite. If there is no prior knowledge, robust estimators can be used as in [32]. The motion parameter estimation algorithm is shown in Figure 6.

5. DYNAMIC EVENT DETECTION AND TRACKING

After motion compensation, the features on the ground plane would be aligned between the two frames, whereas those due to stationary and moving objects would be misaligned. Image difference between the frames would therefore enhance the objects, and suppress the road features. However, the image difference depends on residual motion as well as the

Figure 6: Hierarchical motion estimation algorithm.

spatial gradients at that point. In highly textured regions, the image difference would be large even for small residual motion, and in less textured regions, the image difference would be small even for large residual motion. To compensate this effect, normalized frame difference [40] is used. This is given at each pixel by:

$$\frac{\sum g_t \sqrt{g_u^2 + g_v^2}}{k + \sum (g_u^2 + g_v^2)}$$

where g_u, g_v are spatial gradients, g_t is the temporal gradient. Constant k is used to suppress the effect of noise in highly uniform regions. The summation is performed over a $K \times K$ neighborhood of each pixel. In fact, the normalized difference is a smoothed version of the normal optical flow, and hence depends on the amount of motion near the point. Blobs corresponding to object features are obtained using morphological operations.

Nearby blobs are clustered into one, and the cluster centroids are tracked from frame to frame by an algorithm similar to [17]. For tracking, a list containing the frame number, unique ID, position, and velocity of each track is maintained. The list is empty in the beginning. The following steps are performed to associate the tracks with features:

- At each frame, associate each existing tracks with the nearest feature in a neighborhood window around the track position. Use Kalman filter [2] to update the track with the feature. If no feature is found in the neighborhood window, only time update is performed.

- For features not having a track in their neighborhood, create a new track out of the feature and update it in the next frame.
- To keep the number of tracks within bound, delete the weakest tracks when the number of tracks get too large.
- Merge the tracks that are very close to each other and have nearly the same velocity and are therefore assumed to be from same object.
- Display tracks which have sustained for a stipulated number of frames along with the track history.

For each track that survives over a minimum number of frames, the original ODVS image is used to generate a perspective view [23] of the event around the center of the bounding box.

6. EXPERIMENTAL VALIDATION AND RESULTS

A series of experimental trials were conducted to systematically evaluate the capabilities and performance of the mobile sentry system for event detection and tracking. Three different types of ODVS mountings were utilized to examine the generality and functionalities of the mobile sentry system. The first trial involved a camera on an electric vehicle, the second was using a mobile robot, and the third involved a walking person with a helmet mounted ODVS. The first experiment was done most systematically in order to evaluate the performance. The other experiments are currently in exploratory stage and more work is required to characterize their performance. A related application applying a similar approach to an automobile mounted ODVS is also shown.

6.1 Camera on Electric Cart

The first experimental trial of the mobile sentry utilized the ODVS camera mounted on an electric cart as shown in Figure 7. The cart was driven on a campus road at speeds between 2 and 7 miles/hour (approx. 1 to 3 m/s). The approximate speed of the cart was determined using GPS and used as a-priori motion estimate. It was also observed that the ellipse corresponding to the entire FOV of the ODVS was oscillating, possibly due to relative vibrations between the camera and the mirror, or the automatic motion stabilization in the camera. These were suppressed by estimating the center of the FOV ellipse using a Hough transform, and translating it to a fixed position. The effect of the remaining vibrations were suppressed using the parametric motion estimation process.

Figure 8 (a) shows an image from the ODVS video sequence. The estimated parametric motion is shown using red arrows. Figure 8 (b) shows the classification of points into inliers (gray), outliers (white), and unused (black) points. It should be noted that the outliers are usually identified when the edges are perpendicular to the motion. When an edge is parallel to the motion, aperture problem makes it difficult to identify it. The estimation is done only using the inlier points. Image with the normalized frame difference between the motion compensated frames is shown in Figure 8 (c).

Figure 7: An ODVS camera mounted on an electric cart for a mobile sentry experimental run.

It is seen that the independently moving car and person stand out whereas the stationary features on ground are attenuated in spite of ego-motion. Figure 8 (d) shows the bounding boxes around the moving car and person after post-processing using morphological operations.

Since the algorithm uses a planar motion model, stationary objects above the ground induce motion parallax, and are detected if they are sufficiently close to the camera, and included in the region of interest. Figure 9 shows the detection of a stationary structure as well as a moving person. Only the parts within the region of interest are detected.

The centroids of the detected bounding boxes were tracked over time and the tracks that survived over 10 or more frames were identified. Typical snapshots from these tracks were taken, and the distortion due to ODVS was corrected to get the perspective view looking towards the track position as in [23]. Figure 10 show the snapshots from these tracks, detecting the events.

To evaluate the algorithm performance, the detection results were compared with ground truth obtained by manually observing the video sequence. The performance was compared for two different thresholds on the number of frames in which a track has to survive to be detected as an event. Table 2 shows the detection rate in terms of total number of events (ground truth), and the number of events actually detected. Note that stationary obstacles and shadows are classified as “false alarms”, since they are currently not separated from independently moving objects. Lower threshold increases detection rate, but also increases false alarms. It was observed that 2 events corresponding to a moving person and cart were not detected at all due to the following reasons. The person and cart were quite far and especially the person had a small size in the image. Furthermore, the camera vehicle was turning, inducing considerable rotational ego-motion. Also, the objects were near the boundary of the region of interest that was analyzed. An image of this person is shown in Figure 11.

(a)

(b)

(c)

(d)

Figure 8: Detection of moving objects from an ODVS mounted on an electric cart: (a) Estimated parametric motion of ground plane. Parts of the image corresponding to the cart as well as distant objects are excluded from motion estimation process. (b) Features used for estimation. Gray features are inliers, and white features are outliers. (c) Motion compensated difference image (d) Post-processed image showing detection of the moving car and person. The angle made with X-axis in degrees is also shown.

(a)

(b)

(c)

(d)

Figure 9: Detection of stationary structure in addition to a moving person. Note that only the part of the structure within the region of interest is detected.

Figure 10: Captured events with their IDs (a) Detected persons and vehicles. (b) Shadows and stationary obstacles currently considered false alarms.

Table 2: Performance evaluation. The right column shows the ground truth number of relevant events in the image sequence. The other columns show the number and percentage of events detected by the system. The last three rows show the number of false alarms due to stationary objects and shadows. Ground truth is not relevant here.

Minimum track length	15 frames	10 frames	Ground truth
Total Events	14 (74%)	17 (89%)	19
– Persons	9 (90%)	9 (90%)	10
– Vehicles	5 (55%)	8 (89%)	9
Total false alarms	3	4	N.A.
– Stationary objects	1	2	N.A.
– Shadows	2	2	N.A.

Figure 11: Original image corresponding to a missed events. The moving person and vehicle were far from the camera and near the boundary of region of interest. Also, the camera vehicle was taking a turn, inducing considerable rotational ego-motion in the image.

Attributes such as the time, duration and position of the events were extracted. The camera position at the event time was extracted from the onboard GPS. Assuming that the point nearest to the camera lies on the ground, the event location with reference to the camera could be computed. These were added to the camera position coordinates to obtain the event position. the approximate positions of the camera as well as the actual event were mapped as shown in in Figure 12. Table 3 summarizes some of the events and their attributes.

6.2 Camera on a Mobile Robot

The second experimental configuration was a robotic platform designed in our laboratory. This platform is called Mobile Interactive Avatar (MIA) [19] in which cameras, displays can be mounted on semi-autonomous robot as shown in Figure 13 to interact with people at a distance. The robot was driven around the corridor of our building with people walking around it. Figure 14 shows the detection of moving persons in one of the frames. Snapshots of detected people shown in Figure 15. However, it was noted that the speed of the robot was much smaller than that of people, which would mean that simpler methods could also yield good results in this scenario. Also, the height of the robot was small, and people’s faces could not be effectively captured.

6.3 Helmet mounted camera

The third experimental study for mobile sentry involved a person walking with an ODVS mounted on a helmet as shown in Figure 16. In this configuration, the camera height was approximately 2 meters, which enabled easy capture of people’s faces. The speed of the person with helmet was comparable to that of other people. There was also a considerable camera rotation due to head and body movement, which helped to test the algorithm in presence of large rotational ego-motions. Figure 17 shows the detection of a moving person in one of the frames. To reduce estimation errors, the region of interest for motion analysis was truncated to remove the camera’s own image, as well as the

Figure 12: Map showing the position of the events in red and the camera position at that time in blue. The ego-vehicle track is marked as yellow line. The event IDs are labeled in black.

Table 3: Summarization of event attributes. The left image shows the original ODVS image at the time of the event, middle image shows the output of detection algorithm, and the right image shows the snapshot of the event corrected for ODVS distortion

Event ID: 65		
Event time (snapshot): 16:01:08.3		
Event duration [seconds]: 3.6		
Event position [meters]: (7.1, -6.6)		
Camera position [meters]: (7.8, -5.3)		
Event ID: 109		
Event time (snapshot): 16:01:40.0		
Event duration [seconds]: 1.9		
Event position [meters]: (53.1, 1.8)		
Camera position [meters]: (56.0, -3.1)

Figure 13: Mobile Interactive Avatar: Semi-autonomous robotic system used for evaluating mobile sentry.

objects above the horizon. It is seen that there is significant rotational motion between two frames. In spite of this motion, the moving person is separated from background features such as the lines on the ground. However, if the rotational motion is too large, the detection often deteriorates and tracks get split into parts. Some of the snapshots of detected people shown are in Figure 18.

Unlike the first experiment, the events here consisted of same persons moving around the camera. Also, there was more breaking of tracks due to large rotational motion. Hence, instead of counting the number of events detected, the orientation of tracks in each frame was plotted against time in Figure 19. The identities of the persons were manually recorded, and are color-coded in the figure. The crosses show the track-breaks. It was observed that in a sequence of 5 minutes (3000 frames at 10 frames per second), the persons were tracked when they were sufficiently close to the camera and the rotational motion was not very large.

6.4 Vehicle mounted ODVS

In a related application [15, 22], the event detection approach was applied to an omni camera mounted on an automobile. The actual vehicle speed, obtained from CAN bus was used for initial motion estimate. A video sequence of 36000 frames (20 minutes) was processed and vehicles on both sides of the car were detected by an algorithm similar to the above described one as shown in Figure 20 (a). The distortion due to omni imaging was removed to generate the birds eye view as shown in Figure 20 (b). Figure 20 (c) shows the plots of track positions against time for a segment of the video.

7. SUMMARY AND DISCUSSION

This paper described an approach for event detection using ego-motion compensation from mobile omni-directional (ODVS) cameras. It applied the concept of direct motion estimation using image gradients to ODVS cameras. The motion of the ground was modeled as planar motion, and the features not obeying the motion model were separated as outliers. An iterative estimation framework was used for optimally fusing the motion information in image gradients with a-priori known information about the camera

(a)

(b)

(c)

(d)

Figure 14: Detection of moving persons in an image sequence from a mobile robot. (a) Estimated parametric motion of ground plane. The part of the image in center, which images the camera itself is not used for estimation. (b) Features used for estimation. Gray features are inliers, and white features are outliers. (c) Normalized image difference after motion compensation. The moving person is detected but the lines on the ground are not detected. (d) Postprocessing and tracking output. The track of the detected person with the ID is shown with yellow line.

(b)

Figure 15: Some of the interesting events captured by the mobile robot.

Figure 16: ODVS camera mounted on a helmet. This configuration enables acquisition of snapshots of surrounding people including their faces. However, there is considerable camera rotation due to movement of head and body which should be compensated.

(a)

(b)

(c)

(d)

Figure 17: Detection of a moving person in an image sequence from a helmet mounted camera. (a) Estimated parametric motion of ground plane. There is significant rotational motion that is estimated by the algorithm. (b) Features used for estimation. Gray features are inliers, and white features are outliers. The parts of the image in center, above horizon, and those having small image gradients are not used in estimation. (c) Normalized image difference after motion compensation. The moving person is detected but the lines on the ground are not detected. (d) Postprocessing and tracking output. The track of the detected person with the ID is shown with yellow line. The position coordinates are computed by assuming that the point on the blob nearest to the camera is on ground.

Figure 18: Some of the detected events from the helmet mounted camera.

Figure 19: Time series showing the orientation (theta) of the person with respect to camera. Each color corresponds to a different person. Track breaks are marked by a $\times$

motion and calibration. Coarse to fine motion estimation was used and the motion between the frames was compensated at each iteration. A scheme based on data snooping was used to remove outliers. Experiments were performed by obtaining image sequences from various types of mobile platforms, and detecting events such as moving persons and automobiles, giving satisfactory results.

For future work, we plan to improve the robustness of the system especially for correct localization of large objects. The algorithm currently detects regions containing edges where motion information is significant, but does not respond to uniform areas of large objects. Morphological operations were helpful in combining the detected regions, but a systematic approach based on region-based segmentation and clustering may be more appropriate for getting accurate localization in terms of bounding boxes. It was also observed that the tracks often got broken due to inaccurate localization of the detected blobs. We plan to track entire blobs instead of the centroids to obtain more robust tracks. The events can then be classified into categories such as persons and vehicles using criteria such as size and shape. Learning based approaches such as [34] would also be useful for classification.

The method described above is appropriate for scenes where the background is predominantly planar, and the foreground consists of outliers in form of small objects. If the scene is not that simple, motion segmentation should be performed along with estimation. In case of scenes with multiple stationary planar surfaces, the surfaces have the same parameters for rotation and translation, but different planar normals [14]. Hence, the ego-motion can be parameterized directly in terms of the linear and angular velocity of the camera, and the plane normals of each planar surface. An iterative estimation procedure which estimates each planar surface separately, but uses the estimates it obtains for rotation and translation as starting points for estimating other planar surfaces could make the process more robust to outliers. For example, if the scene consists of features far away from the camera, their ego-motion could be considered almost pure rotation, having only three degrees of freedom. These features could be used to estimate the approximate rotation [40]. This rotation could be used as an initial estimate for the parts of the scene containing ground plane in order to determine the full planar motion. This procedure could be combined with a robust motion segmentation method such as [32] to automatically separate multiple planar surfaces.

Alternatively, the motion parameters can be estimated using a bootstrap method from small patches, and combine the patches having motion consistent with ground plane as done by Ke and Kanade [28]. For 3-D scenes with large variations in depths, structure from motion approach using epipolar constraint [7] is more appropriate. The plane+parallax method proposed by Irani and Anandan [24], can also be used for wide variety of scenes including planar, piecewise planar, and 3-D.

To discriminate between independently moving objects and stationary objects above the ground, the rigidity constraint [24] could be used in the plane+parallax framework. We plan to generalize the piecewise planar motion segmentation as

well as plane+parallax methods for use with ODVS cameras using non-linear motion models for complex scenes and independent motion discrimination.

8. ACKNOWLEDGEMENTS

We are thankful for the grant awarded by the Technical Support Working Group (TSWG) of the US Department of Defense which provided the primary sponsorship of the reported research. We also thank our colleagues from the UCSD Computer Vision and Research Laboratory for their contributions and support.

9. REFERENCES

- [1] O. Achler and M. M. Trivedi. Real-time traffic flow analysis using omnidirectional video network and flatplane transformation. In *Workshop on Intelligent Transportation Systems*, Chicago, IL, 2002.
- [2] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with applications to tracking and navigation*. John Wiley and Sons, 2001.
- [3] R. Benosman and S. B. Kang. *Panoramic Vision: Sensors, Theory, and Applications*. Springer, 2001.
- [4] M. J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding*, 63(1):75–104, 1996.
- [5] T. Boulton, A. Erkin, P. Lewis, R. Michaels, C. Power, C. Qian, and W. Yin. Frame-rate multi-body tracking for surveillance. In *Proc. DARPA Image Understanding Workshop*, 1998.
- [6] S. Carlsson and J. O. Eklundh. Object detection using model-based prediction and motion parallax. In *European Conference on Computer Vision*, pages 297–306, April 1990.
- [7] P. Chang and M. Herbert. Omni-directional structure from motion. In *IEEE Workshop on Omnidirectional Vision*, pages 127–133, Hilton Head Island, June 2000. IEEE Computer Society.
- [8] G. Cielniak, M. Miladinovic, D. Hammarin, L. Goranson, A. Lilienthal, and T. Duckett. Appearance-based tracking of persons with an omnidirectional vision sensor. In *Proc. IEEE Workshop on Omnidirectional Vision*, 2003.
- [9] K. Daniilidis, A. Makadia, and T. Bulow. Image processing in catadioptric planes: Spatiotemporal derivatives and optical flow computation. In *IEEE Workshop on Omnidirectional Vision*, pages 3–12, June 2002.
- [10] G. Danuser and M. Stricker. Parametric model fitting: From inlier characterization to outlier detection. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 20(2):263–280, March 1998.
- [11] K. L. Dean. Smartcams take aim at terrorists. *Wired News*, June 2003.
http://cvrr.ucsd.edu/press/articles/WiredNews_Smartcams.html.

- [12] O. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. The MIT Press, Cambridge, MA, 1993.
- [13] T. Gandhi, S. Devadiga, R. Kasturi, and O. Camps. Detection of obstacles using ego-motion compensation and tracking of significant features. *Image Vision and Computing*, 18(10):805–815, 2000.
- [14] T. Gandhi and R. Kasturi. Application of planar motion segmentation for scene text extraction. In *Proc. International Conference on Pattern Recognition*, volume 1, pages 445–449, 2000.
- [15] T. Gandhi and M. M. Trivedi. Motion based vehicle surround analysis using omni-directional camera. In *Proc. IEEE Intelligent Vehicles Symposium*, 2000. In press.
- [16] T. Gandhi and M. M. Trivedi. Motion analysis of omni-directional video streams for a mobile sentry. In *First ACM International Workshop on Video Surveillance*, pages 49–58, Berkeley, CA, November 2003.
- [17] T. Gandhi, M.-T. Yang, R. Kasturi, O. Camps, L. Coraor, and J. McCandless. Detection of obstacles in the flight path of an aircraft. *IEEE Trans. on Aerospace and Electronic Systems*, 39(1):176–191, January 2003.
- [18] J. Gluckman and S. Nayar. Ego-motion and omnidirectional cameras. In *Proc. of the International Conference on Computer Vision*, pages 999–1005, 1998.
- [19] T. B. Hall and M. M. Trivedi. A novel interactivity environment for integrated intelligent transportation and telematic systems. In *5th International IEEE Conference on Intelligent Transportation Systems*, Singapore, September 2002.
- [20] Haritaoglu, D. Harwood, and L. S. Davis. W4: Real-time surveillance of people and their activities. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 22(8):809–830, Aug. 2000.
- [21] B. Horn and B. Schunck. Determining optical flow. In *DARPA81*, pages 144–156, 1981.
- [22] K. Huang, M. M. Trivedi, and T. Gandhi. Driver’s view and vehicle surround estimation using omnidirectional video stream. In *IEEE Intelligent Vehicles Symposium*, Columbus, OH, June 2003.
- [23] K. C. Huang and M. M. Trivedi. Video arrays for real-time tracking of persons, head and face in an intelligent room. *Machine Vision and Applications*, 14(2):103–111, 2003.
- [24] M. Irani and P. Anandan. A unified approach to moving object detection in 2D and 3D scenes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 20(6):577–589, June 1998.
- [25] B. Jähne, H. Haußecker, and P. Geißler. *Handbook of Computer Vision and Applications*, volume 2, chapter 14, pages 397–422. Academic Press, San Diego, CA, 1999.
- [26] M. Jogan and A. Leonardis. Robust localization using panoramic view-based recognition. In *Proc. Int. Conf. on Pattern Recognition*, pages 136–139, 2000.
- [27] K. Kanatani. *Geometric Computation for Machine Vision*. Oxford University Press, Oxford, 1993.
- [28] Q. Ke and T. Kanade. Transforming camera geometry to a virtual downward-looking camera: Robust ego-motion estimation and ground-layer detection. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages I:390–397, June 2003.
- [29] W. Kruger. Robust real time ground plane motion compensation from a moving vehicle. *Machine Vision and Applications*, 11:203–212, 1999.
- [30] S. Lawson. Yes, you are being watched. PC World, December 27, 2002. <http://www.pcworld.com/news/article/0,aid,108121,00.asp>.
- [31] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *International Joint Conference on Artificial Intelligence*, pages 674–679, 1981.
- [32] J. M. Odobez and P. Bouthemy. Direct incremental model-based image motion segmentation for video analysis. *Signal Processing*, 66:143–145, 1998.
- [33] D. Ramsey. Researchers work with public agencies to enhance super bowl security, February 15 2003. http://www.calit2.net/news/2003/2-4_superbowl.html.
- [34] M. Saptharishi, J. B. Hampshire, and P. K. Khosla. Agent-based moving object correspondence using differential discriminative diagnosis. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, volume II, pages 652–658, June 2000.
- [35] O. Shakernia, R. Vidal, and S. Sastry. Omnidirectional egomotion estimation from back-projection flow. In *IEEE Workshop on Omnidirectional Vision*, June 2003.
- [36] J. Shi and C. Tomasi. Good features to track. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, 1994.
- [37] E. P. Simoncelli. Coarse-to-fine estimation of visual motion. In *Proc. Eighth Workshop on Image and Multidimensional Signal Processing*, pages 128–129, Cannes, France, 1993.
- [38] T. Sogo, H. Ishiguro, and M. Trivedi. N-ocular stereo for real-time human tracking. In R. Benosman and S. Kang, editors, *Panoramic Vision: Sensors, Theory, and Applications*, pages 309–396. Springer, 2001.
- [39] C. Stauffer and W. E. L. Grimson. Adaptive background mixture model for real-time tracking. In *Proc. IEEE Intl. Conf. on Computer Vision and Pattern Recognition*, pages 246–252, 1999.
- [40] E. Trucco and A. Verri. *Computer vision and applications: A guide for students and practitioners*. Prentice Hall, March 1998.

- [41] R. F. Vassallo, J. Santos-Victor, and H. J. Schneebeli. A general approach for egomotion estimation with omnidirectional images. In *IEEE Workshop on Omnidirectional Vision*, pages 97–103, June 2002.
- [42] N. Winters, J. Gaspar, G. Lacey, and J. Santos-Victor. Omni-directional vision for robot navigation. In *Proc. IEEE Workshop on Omnidirectional Vision*, pages 21–28, 2002.
- [43] Z. Zhu, D. Karuppiah, E. Riseman, and A. Hanson. Omni-directional vision for robot navigation. *Robotics and Automation Magazine, Special Issue on Panoramic Robots*. In press.

(a)

(b)

(c)

Figure 20: Detection of moving vehicles in an image sequence using omnidirectional camera mounted on a moving car. The track history of the vehicle over a number of frames is marked. The track id and the coordinates in road plane are also shown. (b) Bird's eye view generated by removing the omni distortion, showing detected vehicles and their coordinates. (c) Plot of the longitudinal position of vehicle tracks against time. The tracks are color coded as red, yellow and green according to increasing lateral distance from the ego-vehicle.