Motion estimation accuracy for visible-light / gamma ray imaging fusion for portable portal monitoring

Thomas P. Karnowski[*a], Mark F. Cunningham[b], James S. Goddard[a], Anil M. Cheriyadat[a], Donald E. Hornback[a], Lorenzo Fabris[a], Ryan A. Kerekes[a], Klaus-Peter Ziock[a], Timothy F. Gee[c]
[a]Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge TN 37831
[b]Lawrence Livermore National Laboratory, Livermore CA 94550
[c]Aldis Inc., Oak Ridge TN 37830

ABSTRACT

The use of radiation sensors as portal monitors is increasing due to heightened concerns over the smuggling of fissile material. Portable systems that can detect significant quantities of fissile material that might be present in vehicular traffic are of particular interest. We have constructed a prototype, rapid-deployment portal gamma-ray imaging monitor that uses machine vision and gamma-ray imaging to monitor multiple lanes of traffic. Vehicles are detected and tracked by using point detection and optical flow methods as implemented in the OpenCV software. Point trajectories are clustered together but imperfections in the detected points and tracks cause errors in the accuracy of the vehicle position estimates. The resulting errors cause a "blurring" effect in the gamma image of the vehicle. In this work, we compare a variety of motion estimation techniques including an estimate using the median of the clustered points, a "best-track" filtering algorithm, and a constant velocity motion estimation model. The accuracy of these methods are contrasted and compared to a manually verified ground-truth measurement by quantifying the root-mean-square differences in the crossing of the gamma ray imaging pixel boundaries compared with a ground-truth system.

Keywords: Radiation monitoring, motion estimation

1. INTRODUCTION

The use of radiation sensors as portal monitors is increasing due to heightened concerns over the smuggling of fissile material. Portable systems that can detect significant quantities of fissile material that might be present in vehicular traffic are of particular interest, especially if they can be rapidly deployed to different locations. To serve this application, we have constructed a prototype, rapid-deployment portal monitor that uses visible-light and gamma-ray imaging to allow simultaneous monitoring of multiple lanes of traffic from the side of a roadway [3][4]. Such a system allows higher spatial selectivity because the gamma-ray imaging instrument can effectively "focus" on a chosen vehicle and generate its radiation signature, avoiding false positives that plague more conventional, simpler systems. In this section we cover the basic principle of operation of the instrument, and then discuss the relevant aspects of the machine vision system.

1.1 Principle of operation

Imaging gamma radiation sources is difficult because the high-energy photons (~50keV to 3MeV) that are given off from radioactive decay do not lend themselves easily to focusing. As an analogy, when visible light is emitted or reflected off an object that we would like to examine, we can place a lens, mirror or systems of lenses and mirrors to focus the photons that come from that object onto a detector. Although there are means of constructing special mirrors to perform focusing of gamma radiation, these are very limited in terms of field-of-view, are difficult to construct, and become ineffective at energies above a few hundred keV. One alternative to image a radiation source is coded aperture imaging **Error! Reference source not found.**[16][17] which uses an indirect method of imaging by means of a blocking radiation sources to cast direction-dependent shadows on position sensitive detectors. In our work we use the Modified

[*] T.P. Karnowski, karnowskitp@ornl.gov, 865-574 5732

Uniformly Redundant Array or MURA mask [1] . These mask patterns have several important properties: first, they have flat side-lobes, which minimizes artifacts; and second, as the number of elements increases, the total area of the "open" holes or white spaces approaches half the total mask area. This is desirable because it allows the collection of more photons. Image reconstruction is conducted by cross-correlating the recorded data with the mask pattern. A one-dimensional imager can distinguish the position of sources in a direction parallel to the imager, but it cannot discern sources above or below the plane of the imager, and only provides minimal information on the range to the source (and then only at close distances). A detailed description of the imaging performance of these systems is beyond the scope of this paper and can be found in the literature **Error! Reference source not found.Error! Reference source not found.**. However, we note that the system of Ziock et al [2] uses the motion of a mobile coded aperture imager to estimate the range of sources though precise knowledge of the location of the imager location and the parallax effect.

Our system for this work is very similar to that of the mobile instrument developed by Ziock et al, however, in our case the imager is stationary with the potential sources moving past the front of the imager. In the moving imager, one can map the gamma-ray images to the world through knowledge of the instrument location, however for the stationary system one does not know where the invidual vehicles are located. This makes it impossible to correctly shift and sum the images to integrated the data for a vehicle transit (except for highly strong sources that show up in the gamma-ray images with only very short integration times). To overcome this, the Roadside Tracker uses machine vision methods to identify and track vehicles in visible-light images. The basic operation is illustrated in Figure 1. For illustration purposes each large pixel or gamma-ray imaging pixel (GRIP) is shown as the constant-angle rays originating from the point below the lane. The field of view of the visible camera system overlaps with and is calibrated to that of a one-dimensional gamma-ray imager. The machine vision system detects vehicles as they enter and exit the field of view, and estimates their position in each frame. The position is used to estimate the times when the front of a vehicle crosses the gamma-ray pixel boundaries. Based on this information, the gamma-ray coded-aperture imager "harvests" the gamma-ray data specific to each vehicle, integrating its radiation signature for the entire time that it is in the field of view [3]. In this fashion we are able to generate vehicle-specific radiation signatures and avoid source confusion problems that plague non-imaging radiation detection approaches to the same problem. More detail on the performance and design of the gamma imager in particular and the entire system are found in the literature [3][4].

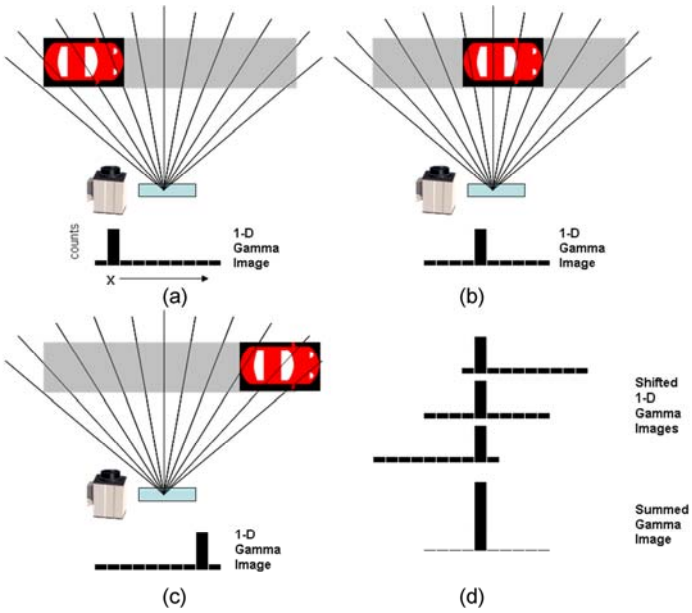

Figure 1. Illustration of visible-light and gamma-ray data fusion. (a) A vehicle drives through the field of view of the gamma imager and machine vision system. The gamma detection is shown as a peak on the left side of the imager. (b) and (c) As the vehicle drives left to right, the peak shifts with the vehicle. (d) Since the machine vision system locates the vehicle in each frame, the gamma images can be shifted the correct amount and added to produce a stronger signal.

1.2 Machine vision system description

Our focus in this work is the visible light / machine vision system. The system has undergone two major designs. The first was a single-roadside proof-of-concept unit that used a single gamma ray imager and a monocular Firewire camera (640x480 pixels) clocked at 15 frames per second (fps) mounted approximately 8 meters above the ground[3]. The second, current system is designed for operation on multilane roadways at ambient velocities. This includes traffic on multilane freeways at speed traffic (~70 MPH) and up to five lanes of traffic[4]. The second system operates with imagers on both sides of the road (although this is optional, it can also operate in "single roadside" mode with reduced performance due to both visible image occlusion and decreased sensitivity of the gamma imager due to greater distances to the potential sources). Our discussion here focuses on the latest advances, which are the use of a stereo imaging approach.

An overview of the machine vision system for one side of the road is shown in Figure 2. The machine vision system performs data acquisition, point detection and measurement, vehicle segmentation, and vehicle motion estimation in a multi-threaded pipelined architecture. The gamma imager and visible light cameras are synchronized through a common clock triggering mechanism. The clock is derived from a GPS time base to precisely synchronize both sides of the road. The gamma imager operates at faster frame rates than the camera system, with both sub-systems using the pulses from the GPS to form a common trigger source for identifying when vehicles cross GRIP boundaries. The camera images are fed into a multi-threaded software pipeline running on personal computers. The first thread coordinates data formatting and writing images to disk. In the newest system three cameras are used: a reference and stereo camera pair, and a third camera (the alignment camera) which images the roof of the trailer and is used to estimate the motion of the stereo-reference images relative to the gamma imager coordinate origin which is pre-defined within the trailer. All cameras **Error! Reference source not found.** are 1032x779 pixels and are clocked at 20 fps. Images are saved to disk uncompressed in Bayer (8-bit) color format.

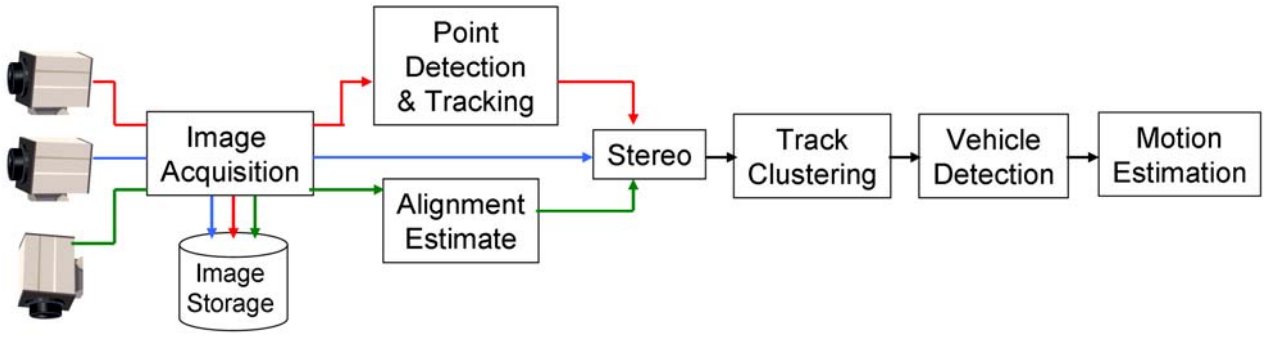

Figure 2. Illustration of machine vision system. A stereo and reference camera pair along with an alignment camera are depicted. After image acquisition and storage, the reference camera (red trace) is used to perform point detection and tracking. In parallel the motion of the camera mast is estimated using the green trace which represents the alignment camera images. The stereo camera is then used along with the point detection and alignment estimate to create 3D estimates of the vehicle points. The final blocks perform track clustering, vehicle detection (identifying when a vehicle has entered and left the field of view), and motion estimation

The second thread in the system pipeline performs point detection and optical flow estimation on the images from the reference camera of the stereo pair using the OpenCV library[5]. A parallel thread uses the images captured from the alignment camera to estimate motion or position changes in the camera system using fiducials on the roof of the trailer. The alignment camera position estimation, the detected points and tracks, and the stereo camera images are then processed in the next thread to perform stereo measurement. The detected points of interest in the reference camera are

searched and found in the stereo camera image. The search is conducted using correlation in an efficient manner: the motion of the point given by the optical flow estimation is used to limit the search area.

The vehicle segmentation relies on the point track clustering. Points are detected[6] and clustered together frame-to-frame using on the Lucas-Kanade algorithm[7]. These form "tracks" which are the basis for the stereo measurements which are projected to the x-y plane (the road) and are clustered together in the track clustering process to form vehicles. Our earlier implementation relied solely on frame-to-frame changes and followed a set of rules inspired by Beymer [8]. Our newest implementation uses "coherent motion regions" and takes sets or "batches" of frames, then clusters the tracks together by solving an optimization process subject to weights attached to the tracks which include relative velocity and spatial location [9]. The entry time and retire time for each vehicle is noted and used to perform motion estimation.

The final step, motion estimation, consists of determining the location of the vehicle in each frame, then performing interpolation between frames to estimate the time of GRIP boundary crossing. This is the format the gamma imaging software receives so it can determine how to co-add the gamma images. Several candidate motion estimation algorithms have been explored and will be discussed in the next section, but as a final note, once the vehicle's position in each frame has been estimated, this is converted to "GRIP crossing events" which are passed via Ethernet links to the gamma imager computers for both the remote and local side.

2. POSITION AND MOTION ESTIMATION METHODS

In our early work[3] we noted that the sensitivity of the gamma imager is such that the optimal peak signal is heavily influenced by two vision-related factors. The first is the estimated range to the vehicle containing a potential radiation source. Our initial assumption that vehicles traveled in the middle of the lane of interest was not correct, and resulted in significant loss of sensitivity. This is rectified in our newer design by better range estimation and the creation of GRIP crossing events at multiple ranges to cover different "slices" within the vehicle. To some extent this use of multiple ranges compensates for motion estimation errors, but the consistency of measurement is important (i.e., consistently reporting the center or edge of the vehicle is more crucial than the exact measurement). The second issue, the primary focus for this paper, is the errors in GRIP crossings caused by motion estimation error. Both errors cause a similar effect, a blurring of the gamma radiation signature and a reduced sensitivity as illustrated in Figure 3. We discuss different options for motion estimation which are directly translated into GRIP values for the gamma imager reconstruction algorithm.

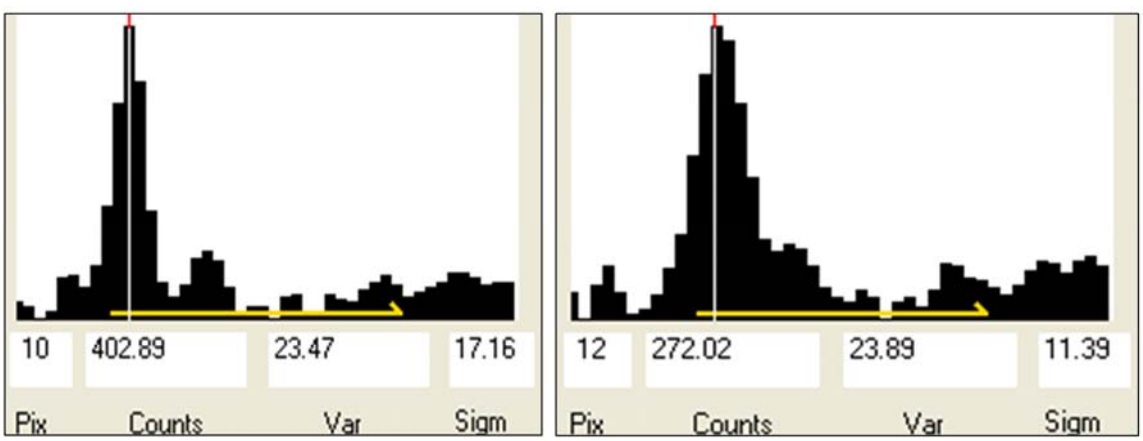

Figure 3. Illustration of reduced sensitivity through blurring effect. These plots show the response of the gamma imaging system to the same vehicle under different GRIP crossing estimates. Left: Image showing good reconstruction. Note the peak value (indicated by "Sigm" on the right side of the plot) is 17.16 and has a sharp shape. Right: Image with slight blur, shown by the widening of the main lobe, and a reduced detection (11.39).

2.1 Position sensitivity

In our system, it is important to have a reliable and consistent estimate of the position of the vehicle in each frame to formulate accurate gamma images. The measurement must be consistent in both the direction of travel (x direction) and in the range (y direction). As an example, in Figure 4 in the top sequence we show the vehicle position estimate centered on the windshield edge in three frames. Since this is consistent, the shifted and summed gamma images focus correctly on the source (as shown in the top right view). In the bottom sequence the middle frame at t_{n+1} has an inconsistent measurement with the other frames. The shifting is based on the position estimate and thus the source is not aligned, and the resulting gamma images are misaligned which results in a degradation of the source estimate. We note that in the long run we also seek to consistently locate the reference point to the same location on all vehicles (i.e., the front bumper). A shift in the reference point between vehicles will create coherent gamma-ray images for each vehicle, but means that source locations identified from the composite gamma-ray image may be shifted along the vehicle. However in this work we seek to determine the best motion estimation method for gamma imaging, not necessarily for visualization.

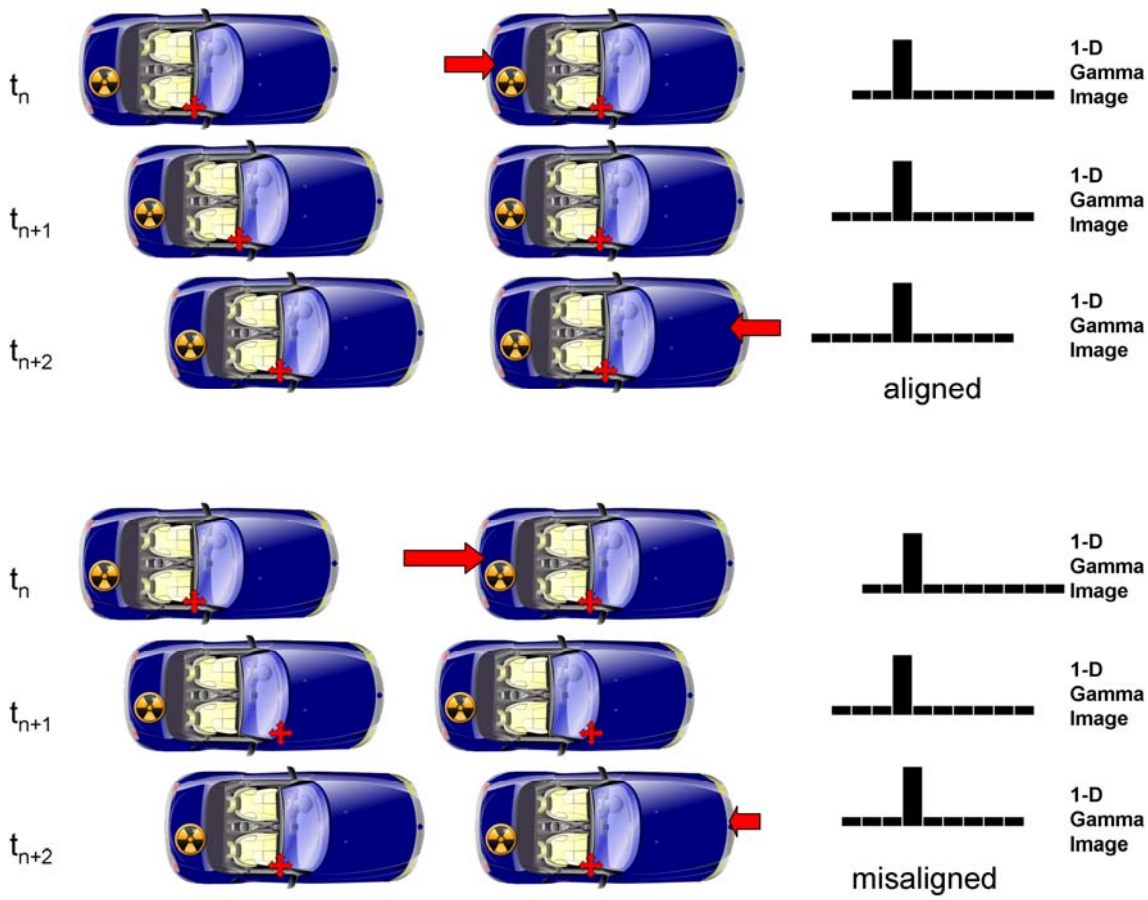

Figure 4. Illustration of x position accuracy requirements. Top left: The machine vision system identifies the edge of the windshield as the vehicle location in 3 subsequent frames. Top center: This results in correct alignment of the frames and (top right) optimal co-adding of the gamma signatures to produce the best possible sensitivity. Bottom: The machine vision system identifies the edge of the windshield as the vehicle location in 2 frames but uses a part of the hood in another, resulting in misalignment of the frames (bottom center) and sub-optimal co-adding of the gamma signatures (bottom right).

In our system it is important to have a reliable and consistent estimate of the range to the vehicle defined as the distance from the gamma-detector plane to the closest approach to a straight-line track past the imager. This is required to formulate accurate gamma images as well as account for lane changes or drift. The gamma imager is sensitive to range and thus we compensate for the fact that we cannot know exactly where the source is within a vehicle by computing the gamma images at multiple ranges or "slices" of the vehicle. However these slices must be at the same range for each frame or the estimate will not be as sensitive as desired. This is illustrated in Figure 5. In the top images the position estimate is again on the windshield for the first frame, but the second frame the estimate has shifted to the door. Consequently to the imager the vehicle has appeared to drift slightly away. As a result the source is no longer at the same range and there is a degradation in the source sensitivity. Thus the y measurement must be consistent from frame to frame.

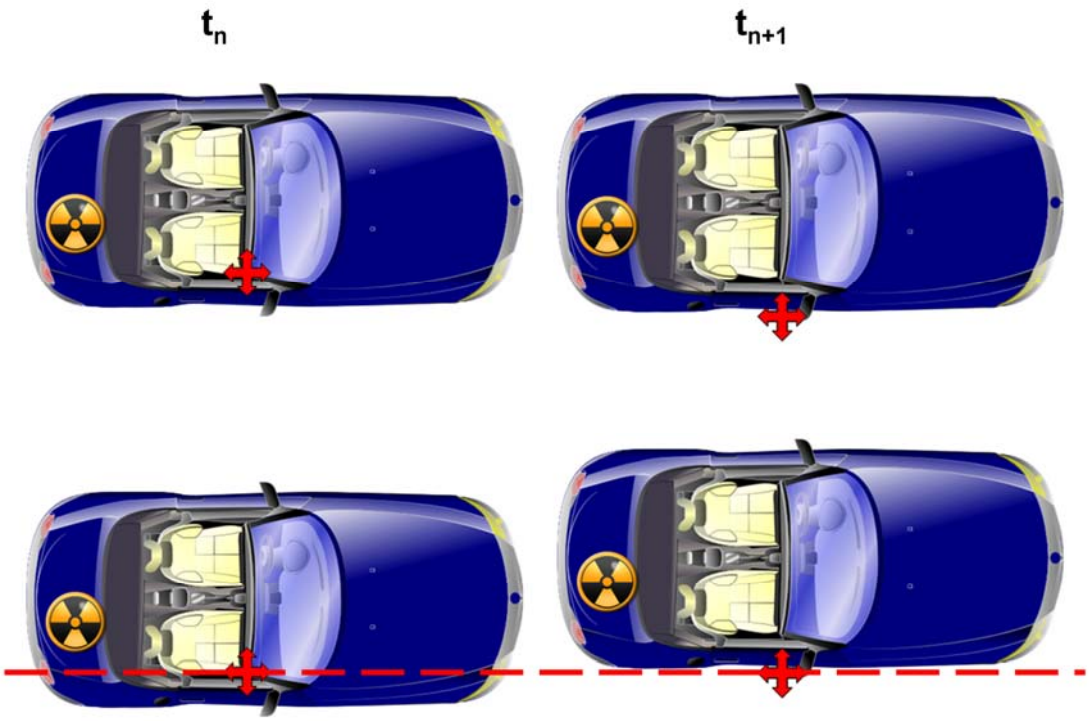

Figure 5. Illustration of y position accuracy. In the top we see that the vehicle location is referenced at the windshield in the first frame and at the door in the second frame. The machine vision system estimates the GRIP boundaries based on the vehicle traveling along the dashed line in the bottom image; thus the effect at the imager is that the vehicle has moved away from the imager when in fact it has not. As a result the source no longer appears at the same calculated range and consequently the sensitivity suffers.

2.2 Manual segmentation and ground truth process

Manual vehicle position and motion estimation is used in this work to constitute a ground truth process. This is shown schematically along with subsequent processing in Figure 6. The images from a vehicle pass are warped to remove lens distortion and projecting to a new view using the system homography transform, where we manually select the two points where each vehicle's tire "touches" the road. Since this operation is performed in the homography transformed view the pixel locations chosen are directly scaled to real-world coordinates. These points are then averaged to estimate the middle of the vehicle. We use test cases where the vehicle does not change lanes and therefore the y position is estimated as the mean of the y positions of the selected points. We then compare these manually obtained results with

the machine segmentation and position estimation of the Roadside Tracker system using the image frames that are common to both the machine vision and ground truth measurements. However the machine vision system is not designed to locate the tires in real-time images, as we do in ground truthing, so we allow for a constant offset in the y location by computing the mean of the y position estimate for each vehicle and offsetting this quantity to the ground truth y estimate (so that the mean of the machine vision detection is identical to the y-mean of the ground truth segment). We also show this in the graph at the bottom left of Figure 6 where machine vision estimate of position (magenta plot) is normalized to the manual estimate (blue plot) by removing a constant offset, but also allows us to evaluate the motion estimation methods for inconsistent y measurements through the vehicle transit which will affect the quality of the gamma signature.

Once these points are selected and specified for each frame, the GRIP crossing times are estimated using linear interpolation of the source location between frames. As mentioned previously, so long as the estimate of the position consistently points to the same feature from frame to frame, our estimates should maximize gamma sensitivity. Since the optimum machine vision tracking point is not based on the tire locations used in the manual approach we expect a fixed offset between the x (and y) locations of the automated and manual data points. As a result, once we have computed the GRIP crossing times, we normalize the crossings to a zero mean to remove these effects. This is illustrated in the bottom right of Figure 6.

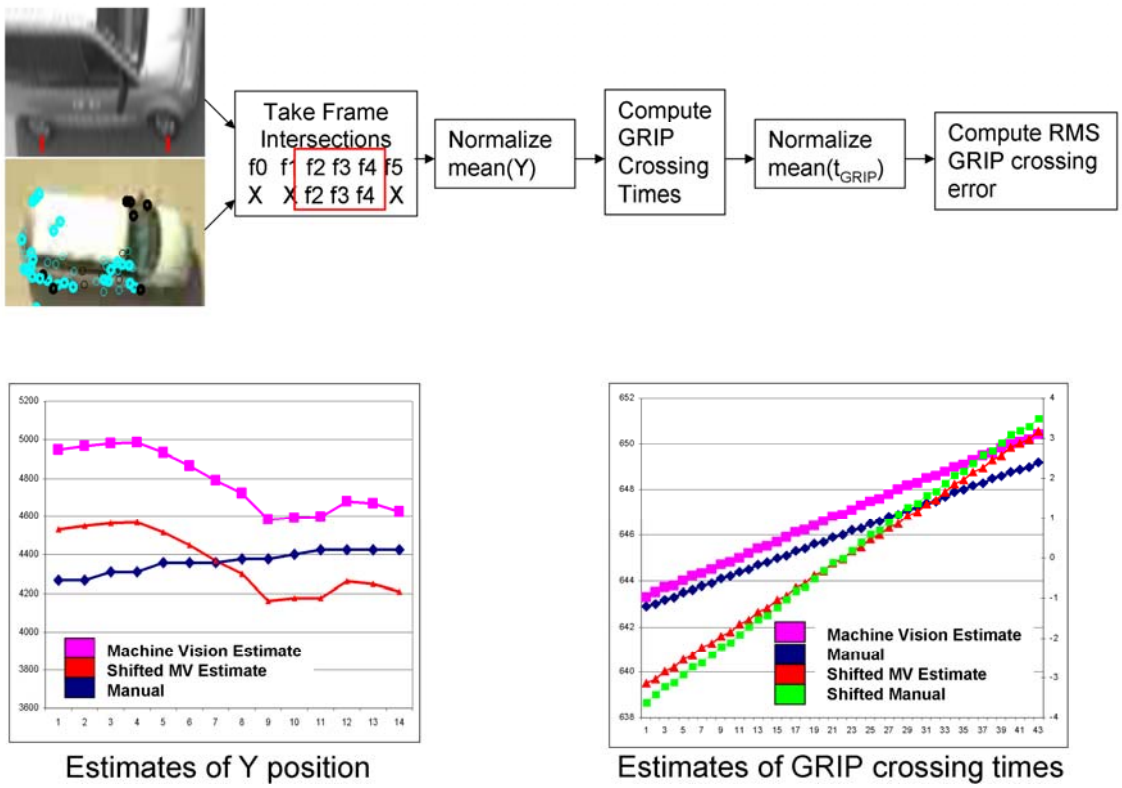

Figure 6. Ground truthing illustration. Top: process flow comparing ground truth values to machine vision system values. Bottom left: plots of y positions before (magenta) and after (red) compensation for systematic differences in Y position estimate as shown in Figure 5. Bottom right: plots of GRIP crossing times before (magenta and blue) and after compensation for systematic differences (red and green) in overall X position accuracy as shown in Figure 4.

2.3 Error sources

Errors in the tracking and measurement process are due to multiple data-dependent sources. In our experience the dominant effects include environment related factors, such as sunlight glinting off the vehicle surface which obscures features on the vehicle, places on the vehicles where features are not as dominant and can be confused from frame to

frame, and non-vehicle points seen in the background, either on the road or on the side of the road. The stereo matching process also produces errors that can result in mismatched points. Some mismatches are detected as errors and thus the system always produces fewer 3D points than corresponding 2D image points. However, there are also errors on the stereo matching process as well where the wrong correspondence point is found and as a result some 3D points are measured inaccurately. A final source of error was shadows on the roadside which result in systematic error in the vehicle position estimate. Some of these effects could be alleviated with other point detection algorithms and other feature similarity methods, but regardless of the method there will likely be errors in determining the true vehicle position.

2.4 Vehicle position methods

A variety of simple estimates of the vehicle position were used for each frame. The first method was used in our previous work and consists of taking the extreme values in x and y for the clustered points in each frame, then estimating the centroid by averaging these extrema. (We will refer to this method "mean of extrema" or MOE for the remainder of this paper.) This is clearly a noisy method but has the advantage of creating bounding boxes that generally cover the entire vehicle. With the increased elevation of the camera system in this work, the noise of the MOE estimate is substantially reduced but is still problematic. In the second method, the mean value of the coordinates of all the points clustered in a frame is computed and used as the position. Finally, the third method is to compute the median value of all the points associated with a vehicle.

2.5 Motion estimation methods

Two different motion estimation methods were used in combination with the vehicle position estimates used above. The first method ("delta") estimates the change in the vehicle position from frame to frame for the entire passage by using either the median change or the mean change. To provide a definite reference point for the vehicle location in each frame, the vehicle location in the approximately middle frame of the transit was chosen and the position in that frame was used as a baseline. The position for previous frames was then estimated by cumulatively subtracting the change in position up to the starting frame, then the process was repeated for frames after the middle frame. The second method was to fit a constant velocity model to the vehicle by taking the estimated positions computed by the MOE, mean, and median vehicle position estimation methods. Combining these two methods with the vehicle position estimates yields a total combination of 8 different motion and position estimation methods for evaluation as shown in Table 1.

Table 1. Motion estimation methods for evaluation.

Name	Position estimation	Motion estimation
Mean of Extrema (MOE)	Maximum and minimum x and y values, averaged	None (uses position estimation method directly)
Mean	Mean of x and y values clustered for a frame	None
Median	Median of x and y values clustered for a frame	None
Mean Delta	Offset from mean value of middle reference frame	Mean value of changes in x and y coordinate in each frame
Median Delta	Offset from mean value of middle reference frame	Median value of changes in x and y coordinate in each frame
Linear Fit to Extrema	Mean of Extrema	Least-squares fit assuming constant velocity
Linear Fit to Mean	Mean	Least-squares fit assuming constant velocity
Linear Fit to Median	Median	Least-squares fit assuming constant velocity

2.6 Best track selection

In our earlier work we learned that the MOE method produced substantial errors in the GRIP crossing time estimates. We found errors on the order of several frames corresponding to temporal errors on the order of 100s of milliseconds. We found that manually selecting a track which "looked right" and using this as the estimate produced a more accurate estimate of position relative to ground truth. To solve this problem, we experimented with a "best track" algorithm which takes all the tracks that are clustered for a particular vehicle and estimates the position using these tracks. The tracks are filtered based on two criteria: the length of track and the goodness of fit to a polynomial model. Thresholds were heuristically established for these parameters and when runs failed to meet the criteria the constant-velocity model was employed as a substitute. For this work, we implemented a very similar algorithm on the stereo motion tracks. Since we studied a variety of vehicle speeds, we choose the longest 25% of the tracks (in time duration) for a vehicle to establish the length-of-track criteria. Each track was then fit to a second-order polynomial and the error computed for each fit. The second order polynomial fit was motivated by allowing for slow changes in the vehicle position, but certainly other models could be used as well. The standard deviation of the error for all the tracks was computed. Then tracks with error greater than one standard deviation were discarded. A final criteria was added with the advent of stereo imaging: tracks with a mean Z-coordinate of 500 mm above ground or lower were also discarded, with the idea that these tracks are likely shadow regions. An example of the Z estimate and the tracks before-and-after filtering is shown in Figure 7. In this example, some of the retained tracks still contain obvious errors. This implies that additional criterion should be investigated to further refine track selection.

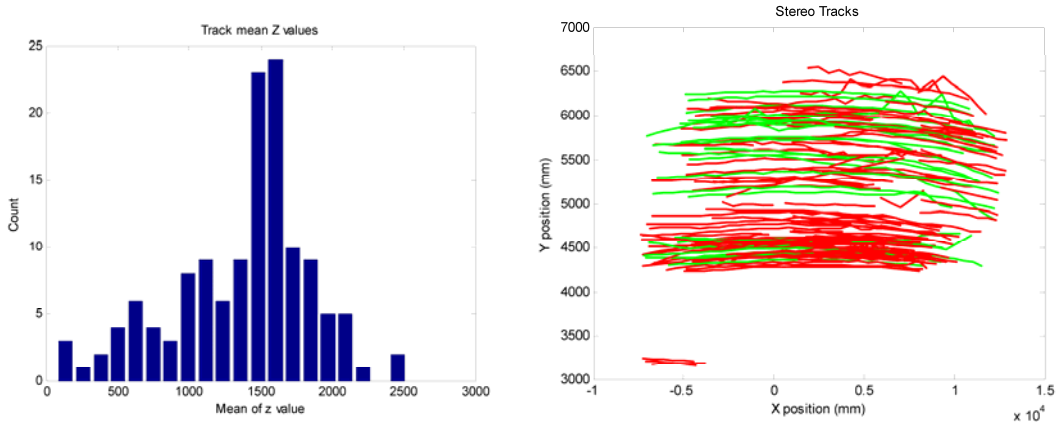

Figure 7. Left: Histogram of z values corresponding to height in mm for tracks for a sample vehicle pass. The vehicle was traveling at 25 MPH. Right: Tracks for the sample vehicle. Red tracks were removed due to size of track, average z value exceeding 500 mm, and/or error fit to second-order polynomial exceeding one standard deviation of the set.

3. EXPERIMENTAL RESULTS

For experiments, twenty different test runs were performed over different velocities in 5 MPH steps from 5 to 55 MPH with two test vehicle passes in each run (one Northbound or left-right and one Southbound or right-left) except two which had only one test vehicle pass for 38 total vehicle detections. We present the results for two basic estimates. First, the standard deviation of the y measurement for each method is examined. As mentioned this gives us an indication of how consistently each method determines the range to the vehicle. The standard deviation of the y measurement was averaged for each method across all runs and is shown in Figure 8. All averages are less than 250 mm and most are around 100 mm. The delta methods show the best performance with a standard deviation under 50 mm. Filtering the tracks actually causes an increase in the standard deviation for some methods, but this may be an artifact of the criteria used in the filtering process.

The second results are the x-location (or the GRIP crossing times). We first took the average of the errors in the GRIP crossing error time for all runs. These are shown in Figure 9. Based on this measurement, the best methods appear to be

the “delta” techniques, the mean delta and median delta, which both perform well regardless of filtering with RMS error below 0.3 frames (or 15 ms with a 20 fps frame rate). For our default method (the linear fit to extrema) the filtered performance is reasonable although it was not one of the top methods. The linear fit models do not work as well as the “raw” or unfit data, which may indicate that the vehicles were not moving at a constant velocity.

In our next set of results, we compared the median delta, mean, and linear fit to extrema methods as example cases. These were chosen because the median delta was the top performer in the average, the mean makes intuitive sense as a good choice and is easy to implement, and the linear fit was our default method. In these results we compare runs as a function of velocity and range (traffic lane) to the vehicle. The results are shown in Figures 10 and 11. We first note that as the velocity increases, the error goes down. Although this seems a little counter-intuitive, we must note that the estimates are based on the clustered tracks. At higher velocities there are fewer tracks and thus fewer are clustered; these that are clustered are more consistent in their velocities. Furthermore, at higher speeds the constant velocity assumption was sounder and thus the linear fit to extrema method is improved. Of course, if the velocity increased more, at some point the finite frame rate and exposure time will be an issue (and blurring of the image itself will occur and degrade performance), but these are not relevant for the velocities used here. As a final note, the performance is worse for the far lane. A possible explanation is that the far lane produces noisier tracks due to interference from the background although this bears further investigation.

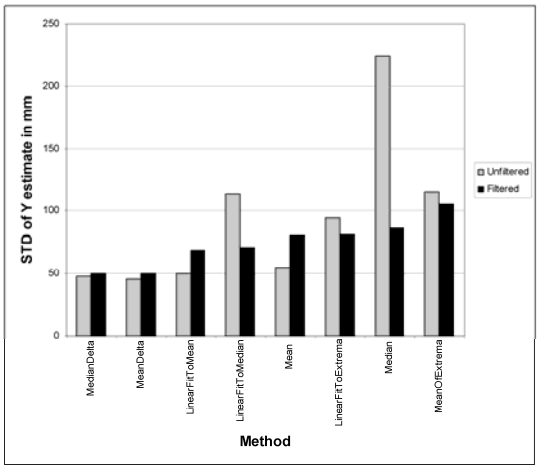

Figure 8. Standard deviation of Y estimate in mm for all methods studied.

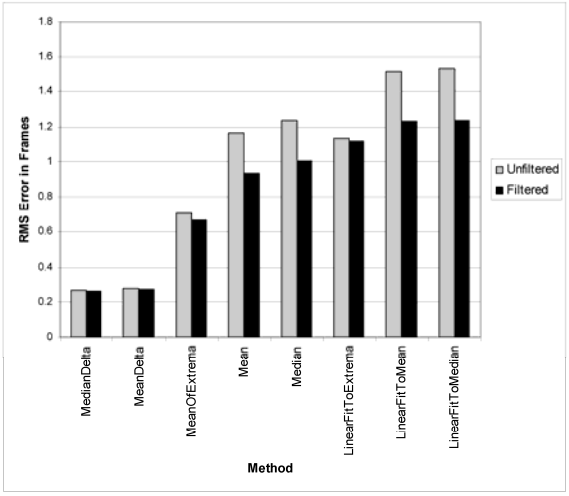

Figure 9. RMS error of GRIP crossing in frames for all methods studied.

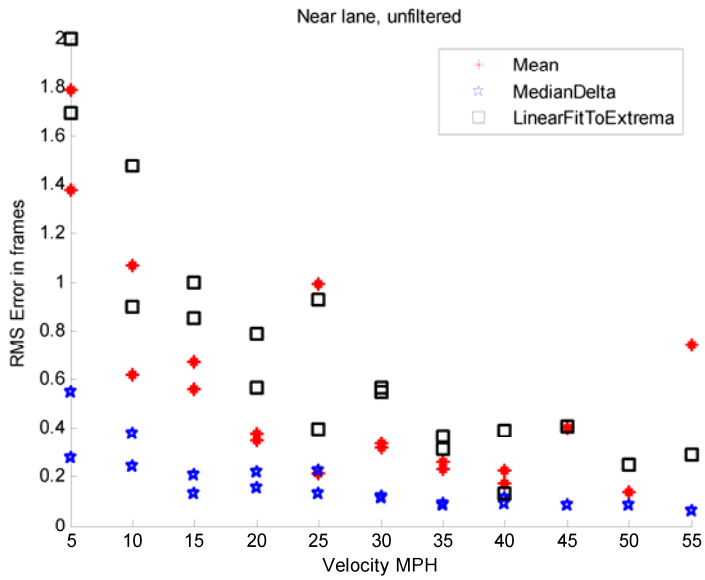

Figure 10. Plot of RMS Error of GRIP crossings for the near lane, unfiltered motion estimates using Mean, MedianDelta, and Linear fit to Extrema method. Performance improves as vehicle velocity increases, although the Median Delta method is the least effected.

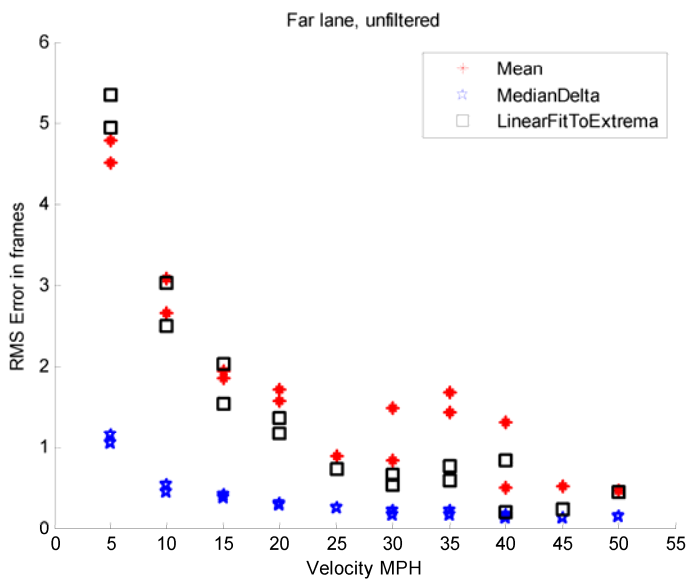

Figure 11. Plot of RMS Error of GRIP crossings for the far lane, unfiltered motion estimates using Mean, MedianDelta, and Linear fit to Extrema method. Performance trends are similar to the near lane case, but overall the error is worse for the far lane.

4. CONCLUSIONS

In this work we examined the performance of different motion estimation methods on clustered tracks of feature points for a vehicle-detection-and-tracking application. In our application the measurement accuracy is important although we are more concerned with consistent performance across all of the frames of a given vehicle pass than we are with the absolute knowledge of the vehicle location. Consistent absolute errors produce only offsets in the source location, while frame-to-frame jitter decreases our ability to detect radiation sources. We found that the median and mean delta methods yield the best performance. They also appear to be relatively robust to the quality of the tracks used, as indicated by their good performance regardless of the best track filtering. For future work we plan to further test the system; quantify the relationship between range, accuracy of tracking, and gamma signature quality; and improve the point detection and tracking by improving the track clustering and point clustering algorithms.

.

5. ACKNOWLEDGMENTS

We would like to thank E.C. Bradley, J. Chesser, and C. Fitzgerald for mechanical engineering and support and also the Department of Homeland Security Domestic Nuclear Detection Office (DNDO) for funding.

REFERENCES

[1] Gottesman, S.R. and Fenimore, E.E., “New family of binary arrays for coded aperture imaging”, Applied Optics, 20(28): 4344:4352 (1989).
[2] Ziock, K.P., et. al., “A fieldable-prototype, large-area, gamma-ray imager for orphan source search”, IEEE Transactions on Nuclear Science 55(6), 3643 - 3653 (2008)
[3] Ziock, K.P., et. al., “The use of gamma-ray imaging to improve portal monitor performance”, IEEE Transactions on Nuclear Science 55(6), 3654-3664 (2008)
[4] Ziock, K.P., et. al., “Performance of the Roadside Tracker Portalless Portal Monitor”, Proc. IEEE Nuclear Sciences, to be published (2009)
[5] http://sourceforge.net/projects/opencvlibrary/
[6] Tomasi, C. and Shi, J., “Good features to track”, Proc. IEEE Conf. on Computer Vision and Pattern Recognition, 593-600 (1994)
[7] Lucas, B., and Kanade, T., “An iterative image registration technique with an application to stereo vision”, Proc. DARPA Image Understanding Workshop, 121–130 (1981)
[8] Beymer, D. et al, “A real-time computer vision system for measuring traffic parameters,” Proc. IEEE Conf. on Computer Vision and Pattern Recognition, 495–501 (1997).
[9] Cheriyadat, A.M., Bhaduri, B.L. and Radke, R.J., “Detecting multiple moving objects in crowded environments with coherent motion regions”, Proc. Sixth IEEE Computer Society Workshop on Perceptual Organization in Computer Vision (POCV), in conjunction with IEEE CVPR (2008)
[10] Barriere, N., et.al., “Soft gamma-ray optics: new Laue lens design and performance estimates”, Proc. SPIE 7437 Optics for EUV, X-Ray, and Gamma-Ray Astronomy IV (2009)
[11] Buis, E.J., et al, “Design aspects of grazing angle multilayer mirrors”, Experimental Astronomy 20(1), 105-113 (2005)
[12] Carolil, E. et al., “Coded aperture imaging in X- and gamma-ray astronomy”, Space Sci. Rev. 45, 349-403 (1987)
[13] Fenimore, E.E., “Coded aperture imaging: predicted performance of uniformly redundant arrays”, Applied Optics 17(22), 3562-3570
[14] Ziock, K.P., et al., “Large area imaging for long-range, passive detection of fissile material”, IEEE Transactions on Nuclear Science 51, 2238-2244 (2004)
[15] http://www.gevicam.com/
[16] R. H. Dicke, "Scatter-Hole Cameras for X-Rays and Gamma Rays", Astrophys. J. 153, L101 (1968).
[17] J. G. Ables, "Fourier Transform Photography: A New Method for X-ray Astronomy", Proc. Astron. Soc. Aust. 4, 172 (1968).