PROJECTILE IMPACT DETECTION AND PERFORMANCE EVALUATION USING MACHINE VISION

Bijan G. Mobasseri
Intelligent Systems Laboratory
Department of Electrical Engineering
Villanova University
Villanova, PA 19085
mobasseri@ucis.vill.edu

Abstract

This paper reports on the development of a machine vision system for assessing targeting accuracy of ballistic, projectile-firing weapon systems[1]*. Current techniques rely on either manual optical sighting or acoustic signature to locate the point of impact. Optical sighting, still the predominant method in many events, is manual and imprecise. Acoustic-based approaches automate the process but require multiple sensor placements. The machine vision system developed here is able to continuously monitor the target, report precise quantitative targeting information and simultaneously provide a color-coded display of impacts. Special provisions have been built-in to account for target plane motion and overlapping impacts phenomenon.*

1. Introduction

Performance evaluation of projectile-firing weapon systems requires quantitative information on impact positions. Conventional techniques have relied on monitoring of the target by a human observer using optical sighting. The observer uses a set of identifying marks on the target, frequently concentric circles, to gauge the performance of the aiming subsystem. Each pair of circles define a constant scoring zone with the score inversely related to the distance of the impact to the target's center. Frequently, scoring zones are identified by numbers from 5 to 10, with the smaller number further away from the center. A typical target after several impacts is shown in Figure 1. The impact position information may then be used in a variety of ways including weapon system retargeting by a machine or a human operator. A more recent approach places a number of sensors around a target, which may be virtual, and acoustic signature is used to infer point of impact information.

[1]Support for this research was provided through a grant from the Ben Franklin Technology Center of Southeastern Pennsylvania and Strainoptic Technologies Inc. grant number 90S.4030R-1.

There are numerous potentials for improvement particularly in extracting quantitative information from the impacts and impact display and visualization. The shortcomings of the existing practices, either at the target design level or in the impact detection and scoring phase, are numerous. First and foremost is accuracy. The basic principal behind the accuracy of a weapon system is the distance, in the Euclidean sense, between the impact centroid and target's center. This basic rule cannot be implemented under the current evaluation practices with any degree of precision. For one, since a scoring zone covers a finite area on the target, it is not possible to discriminate among various impacts placed inside a given zone, i.e. equal scores are given to unequal performances. Conversely, two impacts positioned in neighboring zones, but just on the opposite sides of the boundary, are scored differently and out of proportion to their physical separation. In other words, scoring curve is discontinuous at region boundaries. Another factor affecting accuracy is lack of precise directional information on the impacts relative to the target's center. The standard markings on the target are only able to provide radial position information. Directional information must be inferred by the observer in a non-numeric and inexact terms. Although additional markings can be imprinted to extract such orientation data, this will lead to excessive cluttering of the target plane. The markings on the target are also designed as an aid to the human observer and are not intrinsically relevant. For example, the high score zones are generally in black with the rest in white. This arrangement, as will be seen later, is an impediment, not an aid, to an automated measurement system.

This paper presents the results of a project to develop an automated, precision measurement system for the detection of projectile arrival and point of impact information extraction. The approach selected is through a systematic analysis of digital images captured by a camera trained on the target. The system must be

- Accurate, extracting precise impact position and orientation from full or partial impacts.

0-8186-2840-5/92 $03.00 © 1992 IEEE

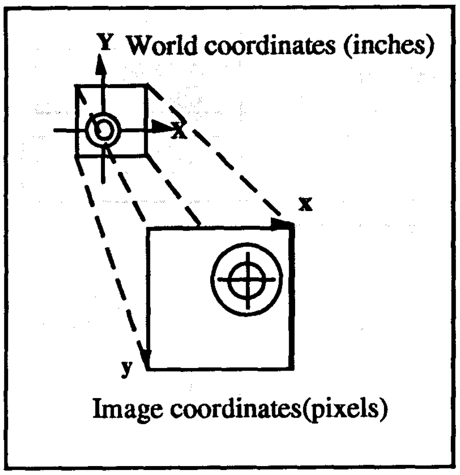

Figure 2. World coordinate and image coordinate differ by the orientation of the axis, scale and reference point.

- Self-organizing, needing little or no operator intervention before or during operation.
- Fast, reporting quantitative data in minimal time while accommodating fast fire rates.
- Robust, needing minimal adaptation to the environment or modification of the environment to fit its needs.
- Low cost for large scale commercial installation.

The rest of this paper is organized as follows. Section 2 describes camera self-calibration process. Section 3 proposes a set of algorithms for impact detection and motion compensation. Section 4 presents the results and experiences encountered in a series of field tests.

2. Self-calibration

From the outset, robustness and autonomy were designated as two primary features of the vision system. These features minimize operator involvement at startup and runtime. Any vision system must undergo a calibration process to establish a proper reference frame. While calibration may be performed manually, the current trend is toward self-calibrating vision systems without requiring active outside involvement. As such, self-calibration has emerged an integral part of active vision (Swain and Stricker, , 1991). As implemented here, operator input is required in only two places in the entire operational cycle, both of which are domain knowledge unavailable to the machine, (1): rough camera alignment and (2): target size or type specification[2].

At the outset, it is important to define the term "self-calibration" in the present context. The location of an impact is characterized by a measurement vector (ρ,θ) in the image coordinate frame where ρ is the length of the vector from the impact centroid to the target's center and θ is its orientation. Therefore, the geometric center of the target in the image coordinate frame is a pivotal parameter upon which all measurements are based. Moreover, operational considerations, as discussed later, require that all image-based measurements be backprojected to the target plane for true length measurement. In achieving these goals, two questions need to be answered. (1): Where is target's center in the *image* coordinate system? and (2): What is the relationship between the image coordinate system and the *world* coordinate system? See Figure 2. Self-calibration here implies that the above two tasks are to be performed by the machine itself. It is of significance to note that this capability is achieved without the addition of any special markings on the target as recognition aid. It was decided that the markings already on paper, the circles themselves, should be sufficient.

2.1 Finding target's centroid

The key parameter affecting all future results is the precision location of the target's center in the image plane. This is a two step process consisting of first locating the boundary of one of the concentric circles followed by extracting the center from it. Among the set of concentric circles, the outer one is the most visible and the least obstructed. Targets used in small arms ranges are generally two tone line drawings. The region closer to the center is frequently the negative, in the photographic sense, of the outside score zones. In either case the circles represent regions of sharp intensity changes hence can be detected easily as points where the intensity gradient exceeds a certain threshold. Automatic thresholding is done using *k-means* clustering. The procedure converges in about 5 iterations and results in a fairly clean segmentation.

The critical phase in self-calibration is the determination of precise target's center in the image coordinate system. The problem is simply that of finding a circle in the image and locating its center. Finding circles in digital images has received considerable attention and many standard techniques are available. The majority of such algorithms are based on some variation of Hough transform. A good literature survey appears in (Yuen et al, 1990). There are also a number of unconventional approaches to the problem (Hanahara and Hiyane, 1990). Our own experience here has shown that

[2]Target type, too, can be automatically sensed. Special tags can be imprinted on the target and then recognized by the software

many established approaches to circle-finding falls short when applied to actual imagery. The standard Hough transform, particularly with unknown circle radius, is too slow for a near real-time performance on personal computers. The algorithm executes in $O(M, N_\theta, N_R)$ where M is the number of edge pixels and (N_θ, N_R) is the granularity of the Hough space. Using gradient direction theoretically limits the search space but in our case the direction angle was not reliable due to the underlying noise. Edge following is another option but never performed reliably, had trouble locating a starting point and got sidetracked later on.

The circle finding technique employed here is a modified version of the chord bisection algorithm (Davies,1987) which is an order of magnitude faster than various versions of Hough transform (Davies, 1990). The principle behind the approach is the simple observation that the center of horizontal and vertical chords of a circle are the loci of the x and y coordinates of the center. In theory, one horizontal and one vertical scan ties down the center coordinates. In a real image, the problem is to assure that what is being looked at is in fact a chord of the desired circle. If the scene is cluttered by objects of arbitrary shapes the center of chords technique as proposed above will most likely fail. The approach suggested by Davies raster scans the image in the horizontal and vertical directions and constructs two 1-D histograms of chord mid-points followed by a peak-finding step. This technique applied to the type of imagery used in this project will fail due to 3 factors, (1): the image consists of a multiplicity of circles not just one, (2): the edge image frequently contains breaks and gaps hence a chord may be ill-defined or non-existent on a particular scan and (3): there are other edge pixels on the image of the target that are unrelated to the circle boundaries. Since it is not a priori known what the identity of individual pixels are on any given scan line, it is difficult to identify a true circle chord. The method developed here uses a 2-D, image-congruent Hough space. Using 2-D cells makes for an approach that is considerably more robust to the effects described above.

To arrive at the estimate of the target's center, the thresholded edge image is raster scanned from top to bottom and then from left to right. For each horizontal and vertical scan, 2 points are extracted, (1): *initial* entrance, and (2): *final* exit. The initial and final entry and exit are significant since on any horizontal, or vertical, scans there are a multiplicity of such points. This is due to (1): concentric circles and (2): isolated noise points. Since we are interested in the location of the outer most circle, all other intercepts are discarded. The estimated initial entrance and final exit coordinates are given by $\hat{x}_e(i)$ and $\hat{x}_x(i)$. The reason for using an estimate is because boundary localization itself is always subject to a finite error. The estimated center for the *ith* horizontal chord is then at $\hat{\mu}_x(i) = (\hat{x}_e(i) + \hat{x}_x(i)/2)$. Similarly, $\hat{\mu}_y(j)$ is defined for the *jth* vertical scan. A two dimensional Hough parameter space, congruent with image space, is formed by partitioning the NxN pixel image into nxn smaller cells and mapping $(\hat{\mu}_i, \hat{\mu}_j)$ into the appropriate cell. The cell showing the highest peak is where the center of the outer circle is most likely to lie. The estimated coordinates of this center is then given by

$$\begin{cases} \hat{x}_c = \dfrac{1}{n'} \displaystyle\sum_{i=1}^{n'} \hat{\mu}_x(i) \\ \\ \hat{y}_c = \dfrac{1}{n'} \displaystyle\sum_{j=1}^{n'} \hat{\mu}_y(j) \end{cases} \tag{1}$$

where $n' \leq n$ is the number of scans falling inside the boundaries of the cell with the equality holding if every row and column is scanned. This approach to center-finding enjoys the inherent robustness of Hough transform by displaying a proportional degradation in performance vs. noise.

2.2 Length measurement

Length measurement in the present context refers to backprojection of distances in the image plane to that in the target plane corrected for imaging geometry and unit conversion factors. In other words, given the impact centroid coordinates measured in pixels, what are the corresponding measurements on the target plane in some physical units?. To arrive at the proper scale factor κ, two pieces of information are required to form a ratio, (1): a known length on the target plane and (2): the corresponding length in pixels in the image plane. The two reference lengths can be extracted from the horizontal and vertical diameters of one of the circles already identified hence additional calibration markings are unnecessary. This ratio, however, is not meaningful if uncorrected for the non-unity aspect ratio η of the frame grabber board. The root of the problem is that length measurements in the image plane are orientation dependent, i.e., identical lengths in the world coordinate system will measure differently, in pixels, in the image coordinate depending on the orientation of the line. There are many approaches reported in the literature to compute the aspect ratio. Due to the presence of concentric circles already on the target, however, it is possible to extract the aspect ratio by simply dividing the vertical and horizontal diameter of the outer circle in the image plane. Both diameters are available from the center-finding stage. This approach is much more straightforward than another technique (Bani-Hashemi, 1991) by not requiring any special markings, patterns or other separate experimental setup. Let (x_c, y_c) be the target center coordinates in pixels and (x_i, y_i) any point on the boundary. Then the radius of the circle in the image plane corrected for the aspect ratio is given by

$$R_i = \sqrt{[\eta(x_i - x_c)]^2 + (y_i - y_c)^2} \tag{2}$$

Since circle boundary pixels are extracted with a finite error, R_i is averaged over (x_i, y_i) contained in the peak cell of the accumulator array. Using actual target diameter information in *mm*, the scale factor is then given by $\kappa = R(pixels)/R(mm)$.

Using the correction factor κ and the target center coordinates in the image plane, the following relationships maps any point (x_i, y_i) in the image to (X_i, Y_i) in the target coordinate frame,

$$\begin{pmatrix} X_i \\ Y_i \end{pmatrix} = \kappa \begin{pmatrix} x_i - x_o \\ y_o - y_i \end{pmatrix}. \tag{3}$$

3. Impact detection

Detection phase is the main component of the vision system. This phase monitors the target continuously and detects the arrival of an impact. Once an impact is detected, target motion is compensated for, actual region of impact is determined and its centroid computed. If a new impact overlaps a previous one, a special algorithm performs centroid calculation on the partially visible impact. Of significance is that the detection algorithm operates asynchronously, i.e. it is not triggered by any external event such as weapon's firing, or any other signal. This is an important property for satisfying the robustness requirement of the vision system. Impact detection can be divided into 3 distinct segments.

3.1 Impact detection and motion compensation

The principle behind impact detection is frame differencing that is corrected for target plane motion and camera noise. The idea is that with everything else fixed, any changes in the image of the target must be due to an impact. To detect an impact, images of the target are captured and successive frames subtracted. Ideally, with no external influences there should be no difference between two successive frames. When a projectile impacts the target the difference image would then contain only the impacted area. Let $f_{t_1}(i,j)$ and $f_{t_2}(i,j)$ be two successive frames. Define a difference image $\Delta f_{t_2}(i,j)$ as follows:

$$\Delta f_{t_2}(i,j) = |f_{t_1}(i,j) - f_{t_2}(i,j)| = \begin{cases} 0 \text{ if } \Delta f_{t_2}(i,j) < T \\ I_{t_2}(i,j) \ \forall (i,j) \ \Delta f_{t_2}(i,j) > T \end{cases} \tag{4}$$

where $I_{t_2}(i,j)$ is the *impact array* at time t_2. The contents of the impact array are simply pointers to those locations in $\Delta f_{t_2}(i,j)$ exceeding the threshold. The centroid of the impact is then given by

$$C = \left(\frac{1}{m} \sum_{i \in I} i, \frac{1}{m} \sum_{j \in I} j \right) \tag{5}$$

where m is the cardinality of the set $I_{t_2}(i,j)$. Computation of (5) completes one cycle of impact detection. In the next iteration, a new impact may have appeared while frame differencing was in progress.. Therefore, subtraction of two consecutively captured images, like at start up, would most certainly return a null image. The reliable measure for change detection is, therefore, the difference between a new frame f_{t_3} and the last one captured prior to subtraction, f_{t_2}. Figure 3 is a timeline diagram showing the chronology of events. At startup, frames 1 and 2 are captured consecutively and compared. The comparison takes Δt seconds during which Impact #1 appears. This impact is detected by differencing frames 3 and 2, taking another Δt, with the process simply iterating from here on. Inspection of Figure 3 reveals a number of constraints and features. First, the two frames being compared are separated in time by Δt. This period may be long enough for the target plane to undergo motions thus invalidating the results of any future differencing operation. Second, there may be more than one impact during Δt. In such a case all impacted areas will be detected but the order will be lost. This observation establishes an upper bound on the firing rate of 1/Δt. Therefore, faster rates can be accommodated if the impact detection cycle time is shortened. Shortening Δt has the added advantage of reducing the chances of target plane shifts during frame subtraction.

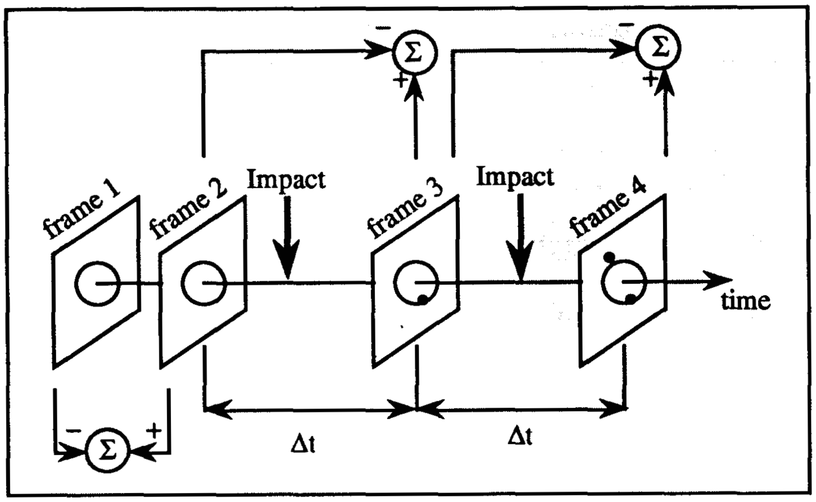

Figure 3. A timeline diagram of impact detection algorithm

The simple approach for impact detection proposed here has several flaws that needs to be addressed. In practice, it was found that there are 2 primary factors, other than impacts themselves, that lead to a finite difference image, (1): camera noise and (2): target plane motion. It appears at first that, with all external parameters fixed, the difference image between two consecutive frames, 1/30 sec. apart, should be zero . A simple experiment with a high quality CCD camera proved that the difference is not zero but in fact a random process (Boie and Cox, 1992). As a result, the presence of a finite difference between two consecutive frames is not necessarily an indication of impact. Since the impacts, too, create an intensity differential, the two sources must be discriminated from each other. The distribution of pixel intensities in the difference image dictates the minimum contrast, about 10, that must exist between the impacted area and its surroundings. This minimum is required to unambiguously differentiate changes due to an impact from those caused by camera noise. It was observed experimentally that pixels change their intensity by at least 20 gray levels following an impact. Therefore, the two events are in fact distinguishable. The above contrast depends greatly upon target background reflectivity or backlighting arrangement and is a crucial parameter in reliable system operation.

The second source giving rise to a finite difference image is target plane motion or deformation as a result of the force of impact. For any frame subtraction to make sense, the two frames must be in perfect registration. Where there is motion, the difference image will also consist of streaks and blobs masking the true impacts. If centroid computation is blindly performed on such an image, the result would be meaningless as it does not reflect the centroid of the impact but that of the entire non-zero pixel set. A technique must be developed that filters the changes due to an impact from those due to the target plane motion. To arrive at this technique, one must examine the characteristic differences between the features arising from the two phenomena. In repeated experiments, it was observed that target plane motion seldom causes more than 1 or 2 pixel interframe disparity. As a result, features in the difference image arising from target plane motion consist primarily of streaks or isolated pixel groups of a few pixels. On the other hand, those features caused by impacts are almost always circular, closed regions with an approximate area of 10 to 15 pixels. It is clear that a thinning operation should be able to eliminate the unwanted streaks and blobs. Conventional thinning operations can be computationally expensive due to the fact that care is taken not to cause breaks or shorten edge ends. In the case here, these concerns are unfounded because, (1): the impacts are closed circular regions hence will be thinned symmetrically and (2): shortening, or eliminating lines or line ends, are of no concern and in fact desirable. These two points allow for a much simpler and faster erosion operation.

To implement the erosion, the impact array $I_t(i,j)$ is used to generate a binary difference image where displaced/impacted pixels are at full and background is at 0 intensity. Within a 3x3 mask, each full intensity pixel is compared with its 8-neighbors. If any pixel in that neighborhood belongs to the background , the center pixel is reset to 0 otherwise it is unchanged. After one pass, all one-pixel wide streaks are eliminated at the same time that impacted regions are thinned, symmetrically, by one pixel at their boundaries. The effect of symmetrical thinning is that impact centroids do not shift from the original unthinned set while other unwanted features are removed. Generally speaking, no more than 2 erosion iterations are needed to correct for target plane shifts. If more violent movements are experienced, erosion will not be effective no matter how many iterations are performed. The reason is that massive target plane shifts, or more seriously bendings, are nonlinear processes. The resulting difference image is not going to be corrected by simple operations unless a full model for such a process is developed. Such modeling, if possible at all, are beyond the realm of performance of affordable vision systems. As it turned out, keeping the target relatively stationary during Δt did not require any special provisions beyond the existing practices.

3.2 Other issue

In many competitive events, impacts tend to be extremely close and at times overlapping. In fact, there could be situation where the projectile simply passes through openings already on the target. The question of overlapping impacts is an interesting problem in computer vision. The problem is basically that of extracting the centroid of a shape from its partial representation. The reliability of such an estimate is based on at least two factors. One is the a priori assumption on the geometry of the impact. Clearly, something must be known about the shape in its complete form. It is not unreasonable to assume a circular, if not a true circle, for the impact area although considerable irregularity exists around the edges. The second factor is the degree of overlap. The larger this factor, the less reliable the centroid estimate will be. In the limit, a vision system alone will not be able to detect a 100% overlap situations. A complete model for centroid extraction from partial impacts has been developed (Mobasseri and Kumar, 1991).

4. Experiments

The vision system described here is implemented on a 33 MHz Intel 486 platform equipped with a Data Translation frame grabber, a Pulnix high resolution CCD camera using a 70-210 zoom lens and a Sony high resolution monitor for image display. This setup has undergone extensive tests at an official range facility without altering any existing operational procedures. Of significance, was the use of standard fluorescent ceiling lights for target illumination with no specialized backlighting. The fixtures, cables, pulleys and all other mechanism used to hold or move the target were left intact. This is significant because such fixtures can easily undergo movements and swings due to impact or air motion. It turned out that the motion compensation algorithm successfully corrected minor translations..

The self-calibration module is started and the operator contains the desired portion of the target, normally the concentric circles, by sliding a resizable rectangular region-of-interest. Once the region is defined, the center-finding algorithm computes the center of the target and identifies it by a cross on the image of the target for inspection. Although it is possible to visually distinguish minute departures from the true center, in repeated experiments the estimate was always within one pixel or less of the actual center. Following this phase, the diameter of the outer circle, in inches or other physical unit, is entered. Rather than using the nominal aspect ratio, we have all the data at this point to measure it experimentally . The measured η is the ratio of the vertical to horizontal diameter of a 5 inch circle in pixels equal to 245/199=1.23 where the published value is 5/4=1.25. Camera setup, a one time operation, and target type are the only instances of operator involvement in the operation of this system.

Following self-calibration, the software enters frame differencing stage of the impact detection phase. The new impact is visually discernible but there are substantial non-zero pixels in the form of streaks and blobs making it difficult To suppress such undesirable effects, the difference image is passed through the erosion operation thereby removing all one or two pixel-wide features. Note that impacts are closed, circular blobs containing for 10 to 15 pixels. The thinning operation, therefore, will not materially affect subsequent centroid computation.

A key concern is the circumstances under which an impact is missed. The most significant factor controlling such an event is lack of sufficient contrast between the impact and its surroundings. This is due to a phenomenon where the projectile generates a perfect cutout upon entrance but an irregular one on the exit side. What happens next is that portions of this irregularity flap back toward the target plane thus restricting available light reflected through the opening. As a result, the contrast between the impact region and its surroundings suffers causing a possible miss. Under such circumstances, the cardinality of the impact pointer array may be too small with several consequences. First, the ensuing erosion operation will reduce the impacted pixels still further making any centroid calculations inaccurate, if not invalid. Second, the overlapping impacts may be completely missed since their size is already at a small number. Although the number of detected impact pixels was at times small, only once in many trials did it result in a miss.

Once an impact is sensed, the impact pointer array is filled with the corresponding pixel coordinates and overlaid on the image of the target. While the rest of the image is displayed in black and white, red is selected to identify the last impact in the sequence. The overlaying of the impact on the target image, as opposed to displaying the difference image, shows the impact location in the proper reference. Quantitatively, the centroid of the impact is extracted and its position vector computed. The information displayed on the monitor would then consist of a numerical score computed from the scoring curve plus centroid bearings relative to the target' center. The latter information is reported in inches above/below or left/right of target's center, Figure 4. Following this phase, the algorithm loops back to the detection phase and the process is iterated. The second impact is similarly detected and overlaid in yellow on the screen already showing the first impact. Similar quantitative information is also displayed for this latest event. It is possible to extend the color encoding into the past to include every impact from the start. Generally speaking, in assessing the performance of the targeting process, the last impact is the most important reference although the capability exists to trace the path of the impacts in a chronologically colored display.

5. Conclusions

The results reported in this work have demonstrated the feasibility of employing a vision system for the performance evaluation of the aiming subsystem of projectile-firing ballistic weapon systems. The strength of such an approach lies in its accuracy, speed, and ease of setup and calibration. The fact that the algorithms have been implemented on the simplest of the platforms available indicates a great potential for speed improvement should the circumstances demand it. Based on our experience, future work should concentrate on means to strengthen the robustness of the algorithms to more substantial target plane motion and impact/surrounding contrast.

Acknowledgments

The author wishes to acknowledge the contributions of M. Senthil Kumar in the implementation and development of the algorithms presented in this paper.

References

Bani-Hashemi A (1991) Finding aspect-ratio of an imaging system. Proceedings IEEE Conference on Computer Vision and Pattern Recognition, June 3-6, 122-126.

Boie R, Cox, I (1992) An analysis of camera noise. IEEE Trans. PAMI 14(6):671-674

Hanahara K, Hiyane M (1990) A circle-detection algorithm simulating wave propagation. Machine Vision and Applications (3):97-111.

Davies ER (1987) A high speed algorithm for circular object detection. Pattern Recognition Letters (6): 323-333.

Davies ER (1990) Machine Vision: Theory, Algorithms, Practicalities. Academic Press, 229-231.

Mobasseri B, Kumar S., An Automated Projectile Impact Detection and Scoring System using Machine Vision Technology, Final Report, 1991.

Swain M, and Stricker M (1991) Promising directions in active vision, NSF Active Vision Workshop, University of Chicago, August 5-7.

Yuen HK, Princen J, Illingworth J, Kittler J (1990) Comparative study of Hough transform methods for circle finding. Image and Vision Computing 8(1):71-77.

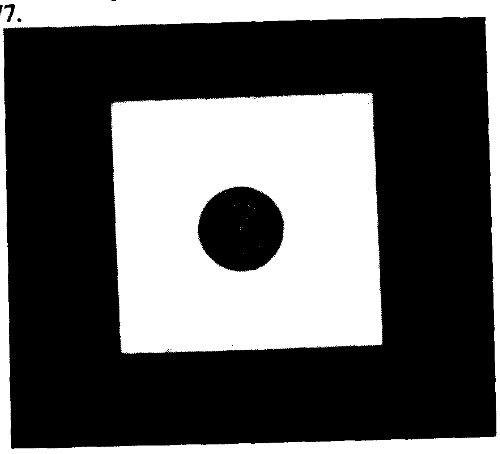

Figure 1. A typical target after a few impacts as viewed through a scope. Impacts are difficult to discern with any degree of precision.

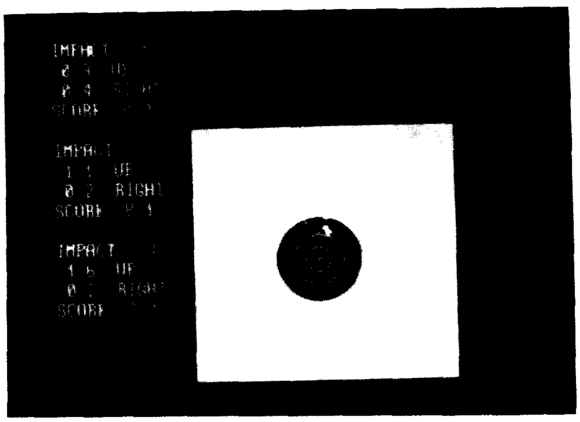

Figure 4. Final output on operator's display. Impacts are color-coded based on their order of arrival. Centroid information, in inches, is displayed on the left relative to target's center.