



New Devices for 3D Pose Estimation: Mantis Eyes, Agam Paintings, Sundials, and Other Space Fiducials

ALFRED M. BRUCKSTEIN* AND ROBERT J. HOLT

Bell Laboratories, Lucent Technologies, Murray Hill, New Jersey 07974, USA

THOMAS S. HUANG

Beckman Institute, University of Illinois, Urbana, Illinois 61801, USA

ARUN N. NETRAVALI

Bell Laboratories, Lucent Technologies, Murray Hill, New Jersey 07974, USA

Abstract. Several unconventional ideas for viewer/camera pose estimation are discussed. The methods proposed so far advocate the use of advanced image processing for identification and precise location of calibration objects in the images acquired, and base pose recovery on the identification of the viewing dependent deformations of these objects. We propose to more fully exploit the freedom in the design of “space fiducials” or calibration objects showing that we can build objects whose images directly encode, in easily identifiable gray-level/color or temporal patterns, the pose of their viewer. We also show how to construct high-precision fiducials, which can determine a viewing direction quite accurately when it is known to lie within a relatively narrow range.

Keywords: Pose determination, space fiducials

1. Introduction

What you see tells you a lot about where you are. People know this and over the years have employed various methods to estimate the pose (i.e. position and orientation) by exploiting images that cameras acquire. Usually it is assumed that in the field of view of the camera there are some feature points (as determined by edges and corners of known objects or on specially designed calibration images), whose relative geometry is a-priori known. If we know the geometry of the feature points and can uniquely identify them in a perspective image taken by a camera whose intrinsic parameters are known, then pose estimation becomes a classical problem of photogrammetry and computer vision. Numerous versions of this problem involving different

numbers of feature points, arranged in various spatial patterns, have already been addressed in the literature. Research was also devoted to pose estimation based on features other than points, like, for example, lines and circles/conic sections or even solid objects such as cubes or planar images displaying various shapes. In all cases, however, pose estimation was addressed by first determining the precise location (via sophisticated image analysis algorithms) of the features whose intrinsic geometry was assumed to be a-priori known, and thereby reducing the problem to the equivalent “algebraic” framework of feature points whose precise perspective projection is provided by the images seen (see, e.g., Basu, 1995; Ganapathy, 1984; Haralick et al., 1994; Holt and Netravali, 1991; Robert, 1996; Tsai, 1987; Bruckstein et al., 1999).

The algebraic framework is very nice indeed and it provides the possibility to do the pose recovery computations that involve solving nonlinear (but polynomial)

*On sabbatical from the Technion, IIT 32000, Haifa, Israel.

systems of equations via iterative, homotopy, or other suitable direct or least square methods. Also it enables one to analyze, via the use of sophisticated techniques from algebraic geometry, the basic limitations and ambiguities involved in pose estimation as a function of the number of features seen and their relative geometry.

Hence, all currently used methods are based on feature identification and require rather extensive image processing and analysis to locate with high enough precision the geometric elements necessary for the (algebraic) computation of spatial pose.

The requirement to do sophisticated image analysis in order to precisely locate features seems to us a very serious drawback of the pose estimation procedures proposed so far. In spite of the fact that it was recognized early on that in many cases of practical importance one can actually produce and deploy in the environment calibration objects and pictures, those objects and images were always designed with the aim of making it easier for the image analysis procedures to precisely locate various features. In this work we approach the issue of pose estimation by more fully exploiting the freedom in designing self-location objects that we may call space fiducials. To do so, we shall not aim at designing objects that enable precise location of some point features, or edges, etc., but rather ask the question: can we build an object that will provide in the image taken by the camera, direct and immediate information about its location and the direction the camera is pointed at. We propose to design objects that, together with the optics and geometry of the camera, implicitly perform reasonably precise analog (optical) computations that result in the camera direction being encoded directly in some gray-level or color or temporal patterns that will appear in the image acquired. In the sequel we analyze several examples of such objects and point out further possible directions for the design of space fiducials that could make 3D pose estimation a more precise and yet a simpler process.

2. Mantis Eyes

The praying mantis is a fascinating insect in many ways. It has, like many other insects, a pair of huge (compared to the body) dome-shaped compound eyes, comprising many thin radial light guides (cylindrical tubes) leading the light towards sensors located at the center of the dome. In images of these insects (that have sufficiently high resolution) one always perceives on the surface of the eye-domes two dark spots (see

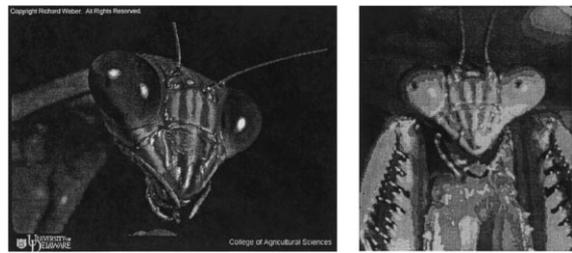


Figure 1. Images of praying mantises. Note the viewpoint dependent dark spots appearing on the eye-domes of the insects.

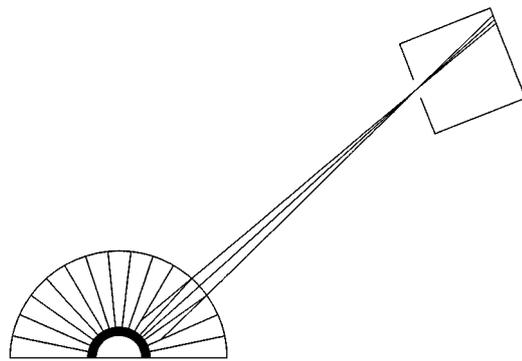


Figure 2. The pupil is seen only through the frustum in line with the optic center.

Fig. 1). If one looks at, or films these insects while moving around, these dark spots seem to follow the viewer/camera, as if the mantis' eye-domes have moving irises with a dark pupil at their centers. Of course, nothing is moving inside the composite eye domes; the apparently mobile dark pupils simply tag the light guides whose directions point at the viewing device's focal point. These are the only "cylinders" that allow glimpses of their deep base while in the others we see only their opaque sidewalls up to various depths (see Fig. 2).

Anyone who ever walked along a tall and wide stack of pipes and tried to see what was on the other side is familiar with this effect. We suggest to use this idea for pose determination in the following way: suppose that a strong light source is placed at the origin and each pipe in a mantis-eye-dome arrangement has a filter that passes light with a different wavelength, or generally a different spectral signature. Then the pupil will change color as it tracks the camera focal point, hence providing, via the RGB values of the image sensor, an immediate and unique labeling of its direction. With two such composite eyes one would then get from the color of the two pupils an immediate readout of the

camera location without any sophisticated image processing. One could even use a color-to-direction lookup table to get immediate direction readouts. (Of course the mantis eye would have to be easily located in the image, and this can be done by having the two domes placed on a readily recognizable brightly colored planar base.)

A moment of thought will show that what we are doing is simply proposing to have a geometric arrangement of light sources and screens that yield a color-coded readout of the direction. One could also achieve this by placing a white light point source at the center of a transparent hemisphere colored in such a way that the color we see uniquely identifies the latitude and longitude of the point illuminated in the image. Clearly a wealth of other similar arrangements can be invented.

3. Agam's Op-Art

The artist Yaacov Agam has built a career around op-art that combines two or more images in dynamic ways based on a simple trick. The two images are decomposed into a series of thin vertical bands of equal size, and these are alternately placed on a serrated surface so that from one side the first image is seen in its entirety and from the other side the second (see Fig. 3). The surface may be regarded as a row of triangular prisms placed side by side.

This idea can obviously be used for pose determination. In fact, suppose that the two images combined in an Agam are an all-black and an all-white image. The all-black/all-white image example shows that we can obtain a direction readout from the average gray-level seen in a region of the picture. The type of pattern we propose for the fiducial is depicted in Fig. 4. While just four triangular prisms are shown in each quadrant of

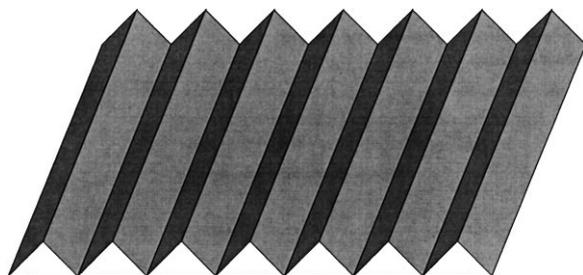


Figure 3. Oblique view of an Agam picture. The plane segments facing up and to the right comprise one image, and the plane segments facing up and to the left form a second image.

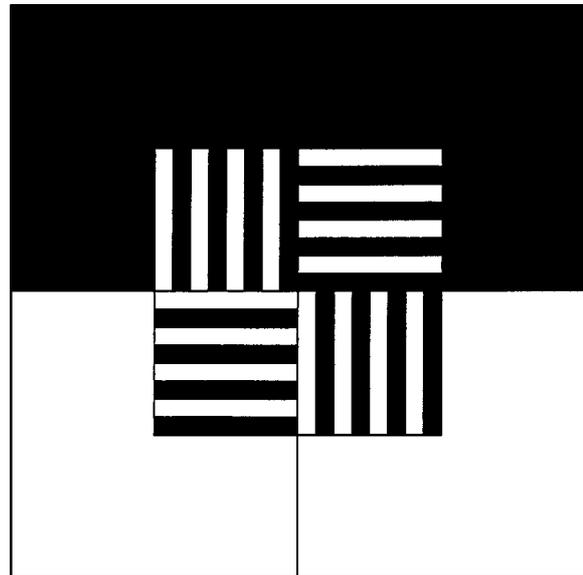


Figure 4. Top view of the proposed space fiducial. The four center squares support rows of triangular prisms, each with one black and one white visible surface. When viewed from a distance, these four squares appear some shades of gray, and the gray-levels observed determine the viewing angles.

this overhead view, we actually intend to have many more, perhaps twenty prisms with rectangular faces of 1×20 millimeters. The patterns of horizontal and vertical strips allow the determination of two angles which suffice to determine the viewing direction.

This fiducial has the remarkable property that if one travels parallel to the plane of its base at a constant rate, the gray-scale value also changes at a constant rate. Consider a piece of the fiducial as shown in Fig. 5.

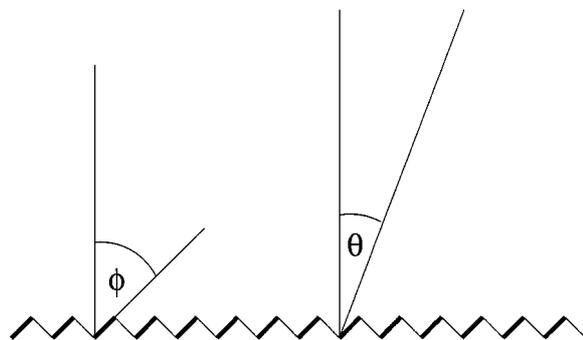


Figure 5. Cross section of the proposed space fiducial. The heavy and light lines indicate surfaces painted black and white, respectively. A vertical line makes an angle ϕ with a plane segment of the fiducial and an angle θ with the line of sight. Orthographic projection is assumed.

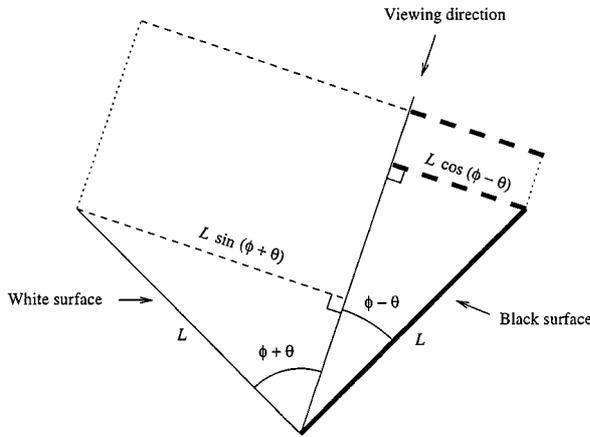


Figure 6. Detailed cross section of one trough. The projections of the black and white surfaces onto a line perpendicular to the viewing direction determine the gray-level seen when they are averaged.

Let the angle a planar segment makes with a line perpendicular to the block be ϕ , and the angle the viewing direction makes with that perpendicular be θ . The angle ϕ may be anything between 0 and 90 degrees, and then the range of θ is $[-\phi, \phi]$. If $\theta < -\phi$ or $\theta > \phi$, we see an all-black or all-white image, respectively. We will assume an orthographic projection since for our purposes the viewing distance will be large in comparison with the size of the fiducial. Also let L be the width of one of the plane segments, and g be the gray-level observed. Now consider the length of the projections of the black and white pieces onto a line perpendicular to the viewing direction as shown in Fig. 6. These lengths are $L \sin(\phi - \theta)$ and $L \sin(\phi + \theta)$, respectively. Therefore the gray-level observed by averaging these two numbers is

$$g = \frac{L \sin(\phi + \theta)}{L \sin(\phi + \theta) + L \sin(\phi - \theta)} N = \frac{N}{2} \left(1 + \frac{\tan \theta}{\tan \phi} \right), \quad (1)$$

where N is the gray-level of white, typically 255. Given a gray-level value of g out of N , with N usually 255, θ is determined from Eq. (1) as

$$\theta = \tan^{-1} \left[\left(\frac{2g}{N} - 1 \right) \tan \phi \right]. \quad (2)$$

The black and white borders in Fig. 4 serve to calibrate the gray-scale. Another possibility is to have a sequence of gray-levels around the perimeter which

could make it easier for the camera to determine the values of the gray-levels in the center of the fiducial.

Note that in the case discussed above we chose to use two uniform images to build the fiducial. However, there are further possibilities. Instead of a fiducial that distinguishes only angles θ in $[-\phi, \phi]$, we can construct one that can distinguish all angles θ in $(-\pi/2, \pi/2)$ by using two non-solid patterns on the alternating strips. Specifically, let x_1 and x_2 be measures of the distances along the width of a “right-facing” and “left-facing” strip, respectively, and assume each strip has unit width. Then the intensity functions $I_1(x_1) = [(1 - x_1)/2]N$ and $I_2(x_2) = [(2 - x_2)/2]N$ provide a pattern in which the observed gray-level increases monotonically as θ ranges from $-\pi/2$ to $\pi/2$. An overhead view of a portion of this fiducial is shown in Fig. 7.

In a manner similar to that as described in Fig. 6, the gray-level value observed from a viewing angle $\theta \in (-\pi/2, -\phi]$ can be shown to be

$$\frac{\int_{-\frac{\sin(\phi-\theta)}{\sin(\phi-\theta)}}^1 I_2(x_2) dx_2}{1 + \frac{\sin(\phi+\theta)}{\sin(\phi-\theta)}} = \frac{\sin(\phi - \theta) \int_{-\frac{\sin(\phi+\theta)}{\sin(\phi-\theta)}}^1 I_2(x_2) dx_2}{2 \sin \phi \cos \theta}.$$

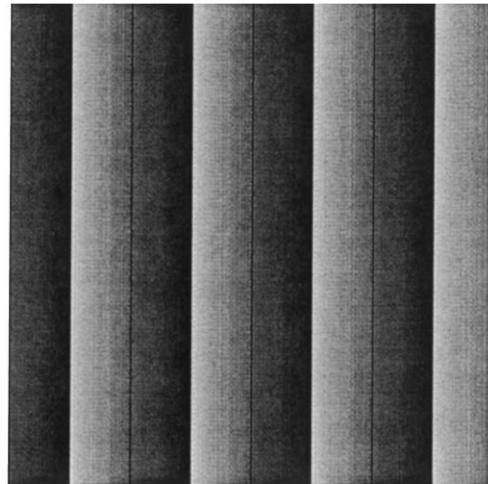


Figure 7. Top and side views of eight strips of the proposed fiducial. The gray-level values of the odd and even strips are given by $I_1(x_1) = 255(1 - x_1)/2$ and $I_2(x_2) = 255(2 - x_2)/2$, respectively. The thin vertical black lines 1/4, 1/2, and 3/4 the way across the image indicate the borders between strips which would not be seen otherwise as the intensity functions are smooth across these borders.

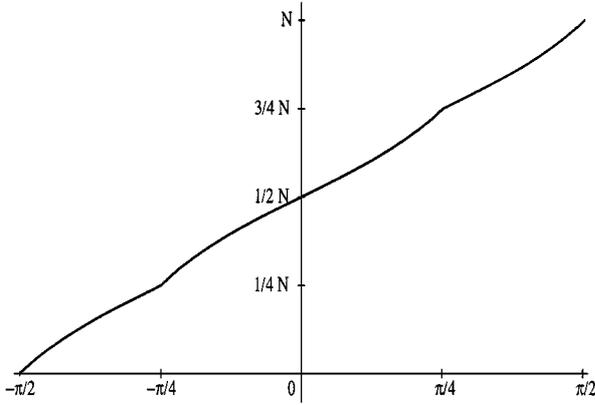


Figure 8. Graph of gray-level value as a function of θ for $\phi = \pi/4$. The function is strictly monotone increasing.

Similarly, the gray-level value observed for $\theta \in [\phi, \pi/2)$ is

$$\frac{\int_0^{1 - \frac{\sin(\theta - \phi)}{\sin(\theta + \phi)}} I_1(x_1) dx_1}{1 - \frac{\sin(\theta - \phi)}{\sin(\theta + \phi)}} = \frac{\sin(\theta + \phi) \int_0^{1 - \frac{\sin(\theta - \phi)}{\sin(\theta + \phi)}} I_1(x_1) dx_1}{2 \sin \phi \cos \theta}.$$

The gray-level seen for $\theta \in [-\phi, \phi]$ is

$$\frac{\sin(\phi + \theta) \int_0^1 I_1(x_1) dx_1 + \sin(\phi - \theta) \int_0^1 I_2(x_2) dx_2}{2 \sin \phi \cos \theta}.$$

Figure 8 shows the plot of the gray-level values as a function of θ for the particular value $\phi = \pi/4$.

Another possibility is to use a surface comprised of several semicircular troughs, as in Fig. 9. This particular surface has the amazing property that if one slice is colored linearly, then the viewing angle is a linear function of the gray-level value observed! This can be seen with the help of Fig. 10. Let the viewing angle be θ , so that $\theta = \angle CDO = \angle OAB = \angle OBA$, and

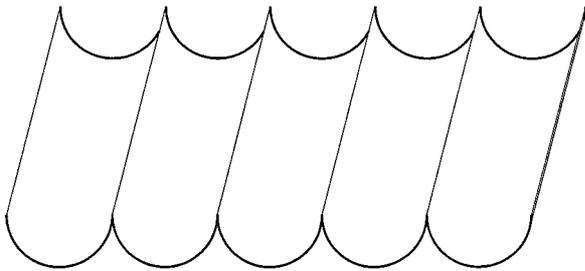


Figure 9. Semicircular fiducial model.

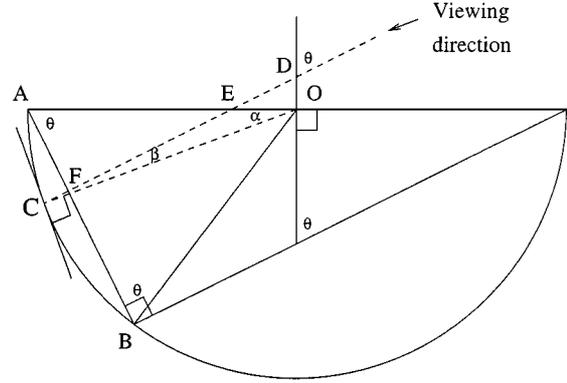


Figure 10. Detail of one slice of the semicircular fiducial.

$\angle AOB = \pi - 2\theta$. Then what is seen in one semicircular trough is the projection of arc \widehat{AB} onto a line perpendicular to the viewing direction, and such a line segment is \overline{AB} . In Fig. 10, α designates the central angle of the semicircle measured counterclockwise starting from A.

The average gray-level observed is found by adding up the contributions from each point on \overline{AB} and dividing by its length, which is $2R \cos \theta$ for a semicircle of radius R . The contribution of a point F on \overline{AB} is the gray-level value of point C , the point on the semicircle in the viewing direction through F , times $\cos \beta$, where $\beta = \angle OCE$ is the angle between the viewing direction and the normal of the semicircle through C . We then have $\angle DEO = \pi/2 - \theta$ and $\angle OEF = \pi/2 + \theta$, and so $\beta = \pi/2 - \theta - \alpha$ and $\cos \beta = \sin(\alpha + \theta)$. The limits of integration are $\alpha = 0$ and $\alpha = \pi - 2\theta$, the measure of $\angle AOB$.

We choose the intensity to be the function $(1 - \alpha/\pi)N$ for $\alpha \in [0, \pi]$. The average gray-level observed is then

$$\begin{aligned} & \frac{N}{2R \cos \theta} \int_0^{\pi - 2\theta} \left(1 - \frac{\alpha}{\pi}\right) \sin(\alpha + \theta) R d\alpha \\ &= \left(\frac{1}{2} + \frac{\theta}{\pi}\right) N. \end{aligned} \quad (3)$$

The above is valid for $\theta \in [0, \pi/2)$. For $\theta \in (-\pi/2, 0]$, a similar diagram and argument shows that the average gray-level value seen is

$$\begin{aligned} & \frac{N}{2R \cos \theta} \int_{\pi + 2\theta}^{\pi} \left(1 - \frac{\alpha}{\pi}\right) \sin(\alpha + \theta) R d\alpha \\ &= \left(\frac{1}{2} + \frac{\theta}{\pi}\right) N. \end{aligned}$$

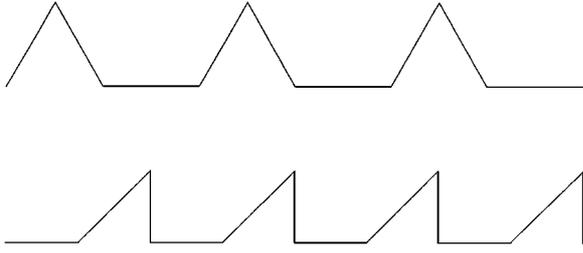


Figure 11. Other possible shapes for Agam-type pictures.

Thus the average gray-level observed for this semicircular fiducial increases linearly from 0 to 1 as θ increases from $-\pi/2$ to $\pi/2$. The graph of the gray-level as a function of θ is similar to that in Fig. 8, with a straight line segment replacing the curve in that figure.

In general, if we have two images $I_1(x, y)$ and $I_2(x, y)$ mapped to a serrated surface we can determine the image seen by looking at it from arbitrary positions in space. We could address an interesting inverse problem in this context. Given what we would like to see, determine whether a pair of functions $I_1(x, y)$, $I_2(x, y)$ exist to achieve this and if these do exist, find them. Also note that other types of serrated profiles might be considered. For example, we might use a surface with planes facing three distinct directions as shown in Fig. 11.

Ideas for Resolution Refinement

A potential shortcoming of the above fiducials is that the accuracy they provide might not be sufficient for the user's purpose. If space permits, we can introduce more fiducials, each of which is geared to a narrow range of viewing directions. This can be accomplished with a series of Agam pictures of the form shown in Fig. 12. The planar surfaces make angles of ϕ_1 and ϕ_2 with the vertical, where ϕ_1 and ϕ_2 are relatively close in value. The planar surfaces alternate black and white in color once more, but this time the whole gray-level spectrum is swept as the viewing angle θ changes from ϕ_1 to ϕ_2 .

As was the case previously, if x measures the distance along the length of the fiducial, then the gray-level value observed increases linearly with x . This can be seen with the help of Fig. 13. The lengths of the projections of the white and black pieces onto a line perpendicular to the viewing direction are $L \sin(\theta - \phi_1)$ and $L \cos \phi_1 \sin(\phi_2 - \theta) / \cos \phi_2$, respectively, where the width of the white strip is taken to be L . Thus the

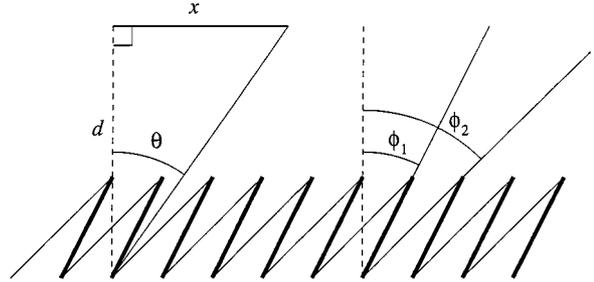


Figure 12. Cross section of the high-precision fiducial. Heavy and light lines represent surfaces painted black and white. A vertical line makes angles of ϕ_1 and ϕ_2 with the white and black surfaces, respectively, and an angle θ with the viewing direction.

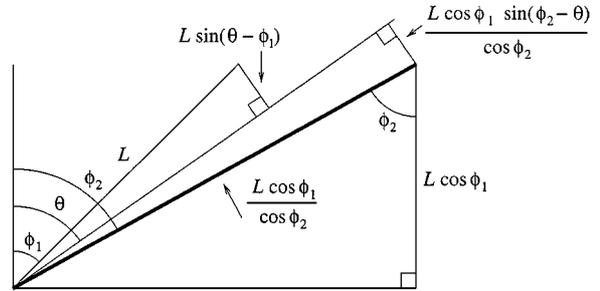


Figure 13. Detailed cross section of one trough of the high-precision fiducial. The projections of the black and white strips onto a line perpendicular to the viewing direction determine the gray-level observed.

gray-level observed by averaging these numbers is

$$\frac{L \sin(\theta - \phi_1)}{L \sin(\theta - \phi_1) + L \frac{\cos \phi_1 \sin(\phi_2 - \theta)}{\cos \phi_2}} N = N \frac{\tan \theta - \tan \phi_1}{\tan \phi_2 - \tan \phi_1}. \quad (4)$$

Equation (4) reduces to Eq. (1) when ϕ_1 and ϕ_2 are replaced by $-\phi_1$ and ϕ_1 , respectively. From (4) we see that the gray-level value increases linearly from 0 to N with $\tan \theta$ as θ increases from ϕ_1 to ϕ_2 .

4. Sundials

A sundial is a device that tells time by measuring the direction toward the sun. We could therefore use the time readout for "sun" pose estimation. Also, a camera could in principle look at a sundial and read out its own position, if it can establish what part of the sundial base is occluded by its needle (see Fig. 14).

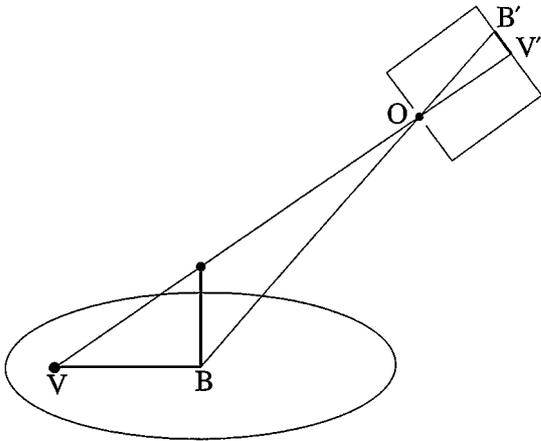


Figure 14. A sundial used as a fiducial. The camera's location is determined by the position of the projection of the needle tip onto the base.

Therefore we could put on the base of a sundial type of fiducial an image that implicitly measures the length of the needle and its angular position. There are, here too, several possibilities. We can combine a planar, circular base with concentric color bands and radial position markers that would provide direct length and angle readouts. This color wheel possibility is shown in Fig. 15.

Many refinements are possible. The base of the sundial can be made concave, and perhaps the most practical shape is a hemisphere (Fig. 16). In this case we obtain a dual to the transparent dome fiducial mentioned in Section 2.

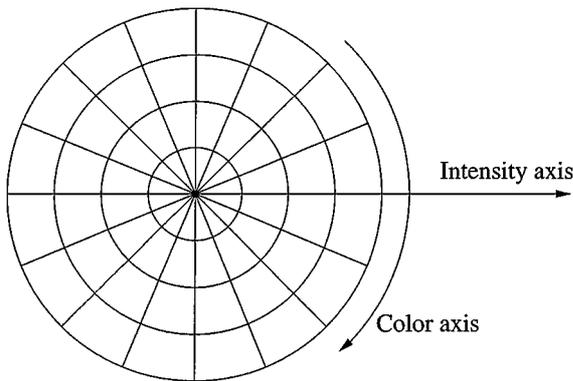


Figure 15. Instead of measuring the length and angle of the needle's projection, we could just read off the camera's position by observing the color where the needle tip's projection lands.

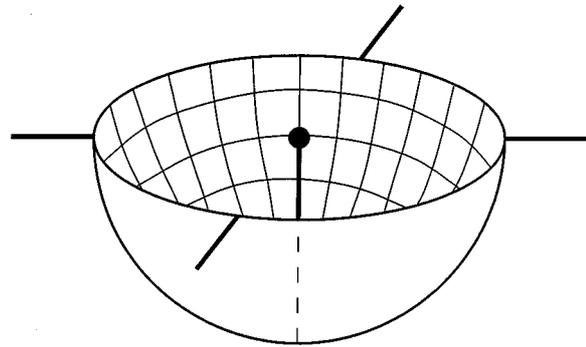


Figure 16. A hemispherical sundial. Latitude and longitude can be directly read from the needle tip's projection.

Note that in many cases we could solve the pose determination problem by "classical" methods of locating some visible fiducial points and then doing calculation; however, these objects provide us by immediate read-out the location up to some precision. Then the location could be refined using more advanced image analysis (locating features to high precision and doing it the classical way).

5. Dynamic Space Fiducials

If we can assume that the cameras we are using take video images (i.e. scan the environment continuously rather than taking snapshots), we can use the temporal dimension too to facilitate the process of pose determination. For example, returning to the Agam idea we can build a dynamic fiducial of the following type: several parallel strips of white material rotate in synchrony about their long axes above a black background (see Fig. 17). The time when the white blades become invisible with respect to the moment when, for example, the background is first completely covered, will tell the camera the direction from the fiducial toward its focal point. Specifically, let θ be the viewing angle as before and ϕ the angle the strips make with the horizontal axis as in Fig. 18. When $\phi = 0$ an all-white view is seen. For $\theta \in [0, \pi/2)$, as ϕ increases the image intensity decreases until $\phi = \pi/2 - \theta$, when the image is all-black. As ϕ continues increasing, the gray-level increases until the image is all-white again, at $\phi = \pi - 2\theta$. This instant is illustrated in Fig. 18. The image is then solid white for $\pi - 2\theta \leq \phi \leq \pi$, and then the cycle repeats. Thus over one cycle, the image is all-white for the fraction $2\theta/\pi$ of the time. We can consider limiting cases of the viewing angle θ . When $\theta = 0$, the image is

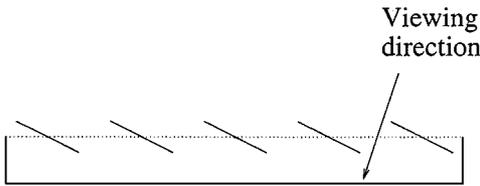
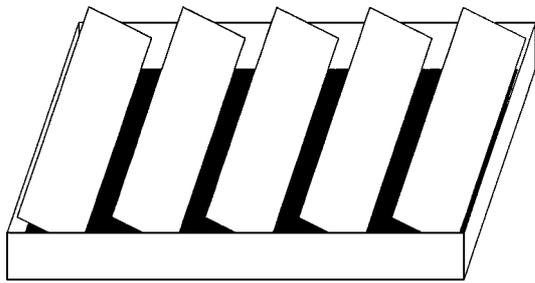


Figure 17. Different gray-level values are revealed as the strips rotate.

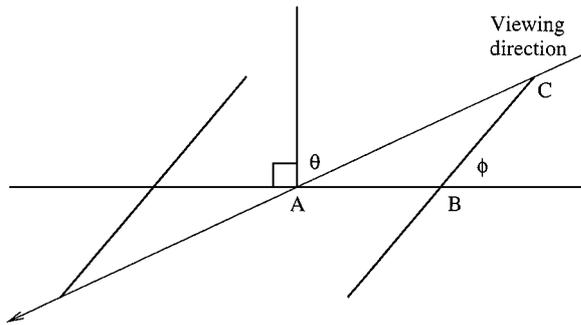


Figure 18. The time instant, for $\theta > 0$, at the beginning of a time interval when the image is all white. $AB = BC$ since the strips just touch when they are horizontal, so $\angle BAC = \phi/2$ and $\phi = \pi - 2\theta$.

all-white at just one time instant (when $\phi = 0$). At the other extreme, as $\theta \rightarrow \pi/2$, the image tends to being all-white the entire time. The plot of the intensity graph over two cycles ($0 \leq \phi \leq 2\pi$) is shown in Fig. 19. An analysis similar to that used in deriving Eq. (1) shows that the observed gray-level is $|\cos(\theta + \phi)| / \cos \theta$ when $0 < \phi < \pi - 2\theta$ and 1 otherwise.

If $\theta \in (-\pi/2, 0]$, the picture is symmetric with respect to the case when θ is positive. When θ is negative, the image is all-white from $\phi = 0$ to $\phi = -2\theta$, and is all-black at $\phi = \pi/2 - \theta$. The image is all-white the fraction $-2\theta/\pi$ of the time. Combining the cases of positive and negative θ , we see that if the image is all-white the fraction k of the time, where $0 \leq k < 1$, then $|\theta| = k\pi/2$. Alternatively, we may observe that the

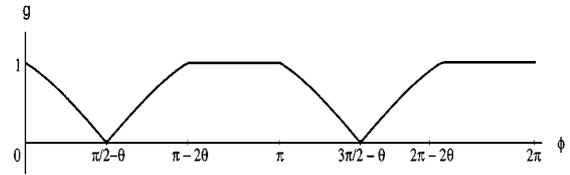


Figure 19. Gray-level values g as a function of time, or of the angle ϕ the strips make with the horizontal axis. The function is $g = \min[|\cos(\theta + \phi)| / \cos \theta, 1]$.

angular difference in ϕ between the “all-black” time instant and the start or end of the “all-white” interval is $\pi/2 - |\theta|$. Therefore if the ratio of this length of time to the length of one cycle (when ϕ increases to π plus its original value) is r , then $|\theta| = (1/2 - r)\pi$. We cannot determine whether θ is positive or negative, or equivalently, if the viewing direction is from the left or the right, from this information alone. This matter can be settled if we have another indication when the white strips are horizontal. If this marker is observed at the start of the “all-white” interval, then θ is positive, while θ is negative if the marker is observed at the end of an “all-white” interval.

As in the Agam fiducial one can use two or more orthogonal arrays to provide location at the intersection of two or several known planes/directions (see Fig. 4).

Clearly this is just one of the many ideas that can exploit the temporal variations seen in some dynamic shapes to extract precise location information. We could for example use temporal, rather than color coding in conjunction with the mantis eye fiducial too.

One could arrange to have LED’s or some other light sources at the base of the hollow tubes that comprise a “mantis eye” fiducial and light these in some ordered fashion that will tell by the delay of seeing the light versus some reference light clock (that could be made always visible) the location of the camera.

6. Concluding Remarks

We have presented several “unconventional” ideas for the design of space fiducials for camera pose determination. The applications areas for those are in camera calibration and self location in unstructured environments like sports coverage and robotic navigation. The main point we are making is that we can and should exploit the freedom in the design of such 3D fiducials to make pose determination a really trivial task. A lot of further research effort should go into this area. Several

interesting issues for further investigation are

- 1) The precision/complexity tradeoffs involved in such fiducial designs. Note that in some applications qualitative information on the camera location and its motion in the environment will suffice, in others very precise data is required.
- 2) The image processing load in the location effort versus the precision obtained. We should design fiducials that minimize the necessity of relying on complex image processing and analysis tasks. In this context previous approaches based on precise location of line/point features are too complicated for real-time video use and the fiducials we proposed could prove to be the only ones capable of solving the problem of continuous self location based on video images acquired.

We have not dealt in this paper with detailed analysis of photometric effects such as interreflections or photometric camera calibration. The issue of interreflections becomes important only if we implement the Agam-type devices with black and white diffusely reflecting surfaces. However, we can also implement them as active devices having their own sources of illumination, for example by having a serrated surface with alternating diffuse transparent and black stripes. In such an implementation the interreflections are negligible and hence we did not address such issues at all.

The calibration is also an important issue that should be addressed in the context of the above discussed location devices. Here too we can actively address the problems via the provision of graded linear or non-linear, but a priori known (in fact, suitably designed), bands of gray-level calibration charts. The topic is by no means a trivial one and should be carefully addressed in the design of each device; however a thorough discussion of it here would be impractical given the scope of our paper, but see e. g. the works (Nayar et al., 1991; Nalwa, 1993; Healy and Kondepudy, 1991; Debevec and Malik, 1997).

The main part of our paper is the observation that we can design a variety of objects exploiting active illumination, color, surface reflectance, and transparency effects, and even controlled motion to solve problems of self location either precisely or qualitatively. In fact optical beacons have long been used in aviation for such purposes. A thorough check for patents revealed two highly relevant ones in this line of thought. In Bergkvist (1979) and some related Swedish patents, Bergkvist describes a clever “device for optic, preferably vi-

sual determination of a certain plane” based on Moiré patterns that arise as two transparent screens with linear gratings, one occluding the other, are viewed from different angles. This is a very nice idea clearly in line with the methods proposed herein. Also Kunkel proposed in his patent titled “Sensing apparatus for determining the relative position between two bodies with plural emitters and a shading member” (Kunkel, 1987) a device based on the occlusion caused by a perpendicular rod on a planar area having a 2-D pattern of sources of light. This idea is clearly an implementation of the sundial fiducial as described herein. The fact that such ideas have already been proposed (as we found out after our paper was submitted), only strengthens our belief that self location and pose determination with active devices as described above are viable alternatives to the methods based heavily on high accuracy image analysis, and research in this direction should be pursued further.

References

- Basu, A. 1995. Active calibration of cameras: Theory and implementation. *IEEE Transactions on Systems, Man, and Cybernetics*, 25:256–265.
- Bergkvist, L.A. 1979. Device for optic, preferably visual determination of a certain plane. U.S. Patent #4,166,699.
- Bruckstein, A.M., Holt, R.J., Huang, T.S., and Netravali, A.N. 1999. Optimum fiducials under weak perspective projection. *International Journal of Computer Vision*, 35:223–244.
- Debevec, P.E. and Malik, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proceedings of SIGGRAPH 97*, pp. 369–378.
- Ganapathy, S. 1984. Decomposition of transformation matrices for robot vision. *Pattern Recognition Letters*, 2:401–412.
- Haralick, R.M., Lee, C.-N., Ottenberg, K., and Nölle, M. 1994. Review and analysis of solutions of the three point perspective pose estimation problem. *International Journal of Computer Vision*, 13:331–356.
- Healy, G. and Kondepudy, R. 1991. Modeling and calibrating CCD cameras for illumination invariant insensitive machine vision. In *SPIE, Vol. 1614: Optics, Illumination, and Image Sensing for Machine Vision VI*, pp. 121–132.
- Holt, R.J. and Netravali, A.N. 1991. Camera calibration problem: Some new results. *CVGIP: Image Understanding*, 54:368–383.
- Kunkel, B. 1987. Sensing apparatus for determining the relative position between two bodies with plural emitters and a shading member. U.S. Patent #4,710,620.
- Nalwa, V. 1993. *A Guided Tour of Computer Vision*. Addison-Wesley.
- Nayar, S.K., Ikeuchi, K., and Kanade, T. 1991. Shape from interreflections. *International Journal of Computer Vision*, 6:173–195.
- Robert, L. 1996. Camera calibration without feature extraction. *Computer Vision and Image Understanding*, 63:314–325.
- Tsai, R.Y. 1987. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics and Automation*, 3:323–344.