5±2 Eigenimages Suffice: An Empirical Investigation of Low-Dimensional Lighting Models

Russell Epstein Peter W. Hallinan Alan L. Yuille
Division of Applied Sciences, Harvard University, Cambridge, MA 02138

Abstract

Recently, Hallinan [1] proposed a low dimensional lighting model for describing the variations in face images due to altering the lighting conditions. It was found that five eigenimages were sufficient to model these variations. This report shows that this model can be extended to other objects, in particular to those with diffuse specularities and shadows. We find that sharp specularities and shadows cannot be well represented by a low dimensional model. However, both effects can be adequately described as residuals to such a model. We can deal with occluders in a similar way. We conclude that low dimensional models, using 5±2 eigenimages, can be usefully extended to represent arbitrary lighting for many different objects. We discuss applications of these results to object recognition.

1 Introduction

In object recognition, the two most common approaches to handling the problem of lighting variations have been (1) extracting lighting invariant features and (2) using a collection of templates or feature sets spanning the entire range of anticipated conditions. Neither approach has proved tenable [1], [2]. A third approach, verified for human faces in [1], [3] is to model all the possible images of an arbitrarily lit face of fixed pose using a small number of parameters. This lighting model can then be augmented to handle geometric distortions and used to recognize faces under general lighting conditions [4]. This paper extends the previous work of [1], [3] from faces to arbitrary objects.

Recently [5] and [6] observed that if an object is Lambertian then an image of it under arbitrary lighting conditions can be obtained by taking a linear combination of three basis images. Though this result has some empirical support [6], it not only excludes reflectance functions having significant specular lobes and spikes (see Nayar et. al.[7]), but also ignores critical effects of shadows, mutual inter-reflections, and occlusions [8]. Any generally applicable approach to modeling arbitrary objects must account for these effects. Our strategy, therefore, is to investigate experimentally the variations of images of objects caused by changes of lighting conditions. We study objects which contain sharp specularities and shadows, and investigate the effect of partial occlusions.

The key idea underlying our lighting model is that any lighting condition can be decomposed into a sum of point light sources. Therefore, all the possible images of an object under different lighting conditions can be represented as linear combinations of images formed from point light sources, assuming viewpoint and object pose remain constant. By taking the principal components of these point light source images, one gets a low-dimensional space in which all lighting conditions can be represented. Unlike other PCA analyses (eg [9]) that have confounded intensity and geometric variations, our approach explicitly isolates lighting variations and guarantees that all significant lighting effects are captured. For faces, [1] found that five eigenvectors sufficed to represent most of the variance in lighting.

Our results demonstrate that, as long as the pose, surface properties (e.g. texture), and surface geometry of an object are held constant, then:

- the first few eigenvectors describe the lambertian component,
- the next few eigenvectors encode the specular lobe,
- the remaining eigenvectors encode specular spikes, small scale cast shadows, and other irregularities such as partial occlusions.

As a result, we conclude that for many objects,

- 5 ± 2 eigenvectors will suffice to model the lambertian and specular lobes,
- specular spikes, small shadows and occluders can be treated as residuals and eliminated by projecting the original image onto the low dimensional eigenvector model,
- the sampling of lighting directions required in the training set increases with both the specularity and the complexity of the surface geometry.

2 Constructing the Lighting Models

The following methodology was used to construct the lighting models for each object. Seventy images were taken of each object under different lighting conditions. The lighting in each image was from a small area source (floodlight) at a distance of about six feet (see [3] for details.) This light could be moved along the surface of a sphere whose center was the object and whose North Pole was directly above the object. The light was moved along the sphere's lines of latitude and longitude in 15 degree intervals such that the

0-8186-7021-5/95 $4.00

lighting direction varied over the entire right front side of the object. These images formed the *dense* training set. The *sparse* training set was a subset of 20 images from the dense set. The lighting for the images of the sparse set varied in 30 degree intervals.

Two eigenimage models were constructed for each object by performing principal component analysis on the sparse and the dense data sets. Additional images of the objects were taken under ambient lighting conditions. These images were used to evaluate the models' ability to reconstruct novel images which were not in the training sets. In addition, when the sparse data alone was used, the remaining images of the dense set were used to test the sparse model.

The quality of the eigenimage models was measured by using a goodness of fit criterion as follows. First define the reconstruction of a preprocessed input image J^i from an eigenimage basis $\{S_{l,b}\}$ as

$$R^i_{k,b} = \sum_{l=1}^{k} < J^i, S_{l,b} > S_{l,b} \quad (1)$$

where k is the number of eigenvectors in the basis and b labels the different bases corresponding to different sample sets. Then the quality of the reconstruction can be measured by the goodness of fit function

$$\epsilon(i,k,b) = 1 - \frac{||R^i_{k,b} - J^i||^2}{||J^i||^2} \quad (2)$$

Since both $R^i_{k,b}$ and J^i have already had the mean image subtracted from them, and since $||\cdot||$ is an L2 norm, this criteria severely penalizes mismatches. It also provides us with a clearcut method to trade model accuracy for fit.

Figure 1 contains graphs showing the mean goodness of fits (i.e. the original images minus the reconstructions) across both training and test data sets for bases constructed from the sparse samples only. These graphs show clearly that, with the exception of highly specular objects, improvements in both average and worst case performance begins to taper off at around 5 eigenvectors. The results for bases constructed from denser sample sets were comparable. (For more details, see [10].)

3 Objects Without Specular Spikes

We start by examining a basketball, an object with a very simple geometry, no self-shadowing, and few specularities. The eigenvectors calculated from the 20 images of the sparse set are shown in Figure 2. Table 1 demonstrates that a very small number of eigenvectors are necessary to account for most of the variance in the data– three eigenvectors are sufficient to account for 94% of the training data while five account for 98%. Thus, the training set data is easily represented by a low dimensional eigenspace.

This eigenspace should be able to represent images taken under all possible lighting conditions, not just those in the training set. To examine this, we make use of images from the dense set. Figure 3 shows the results of reconstructing these images using three eigenvectors. These reconstructions are almost impossible to distinguish from the original set. Three eigenvectors are quite sufficient to represent the lighting for this object.

However, all of the first five eigenvectors for the ball are interesting because they are similar to the eigenvectors found for faces (see [3], [1]). For faces these five eigenvectors formed a stable set, very similar between different people, which corresponded to lighting from the front, side, top/bottom, corners, and extreme side (though not necessarily in this order). The results for the ball suggest that this result holds generally for smoothly curved objects with little specular component.

Table 1 shows that the lower eigenvectors account for more of the variance for the ball than they do for faces. Five eigenvectors only account for 94% of the variance for faces, as opposed to 98% for the ball. Figure 4 explains the difference. On the left is an image not in the training set. The second image is the projection of this image onto the first three eigenvectors. Observe that the specular lobe on the forehead is imperfectly reproduced. The third image is the projection onto the first five eigenvectors; here the specular lobe is more prominent. Thus, the face requires more eigenvectors than the ball largely because the face is more specular. Observe also that the shadow under the nose is also imperfectly reproduced in the projection. This is because the face has more complex geometry than the ball and hence has more shadows. In particular, the shadow under the nose causes problems because it varies significantly in shape and position as the lighting direction changes. So it cannot be well modeled by a linear combination of images, because linear combination does not allow for *geometric* interpolation between the basis images (see discussion in [11].) This problem is even more pronounced with objects having more prominent specularities and shadows. The fourth image is the residual image, obtained by taking the difference between the original image and its projection onto the first five eigenimages. (In other words, the fourth image shows the absolute value of the difference between the first and third images.) The difference between the two images is actually quite small, corresponding to an average over all pixels of 4.687 greyscale levels, with a maximum pixel difference of 75 greyscale levels. (To compare, the pixels of the original image range over 207 greyscale levels, with a standard deviation of 33.4.) The areas of greatest difference are the shadow under the nose, and the hairline, both of which are not well represented by five eigenvectors.

To investigate this further we calculated the eigenvectors for the images of an artificial stuffed parrot. This object has more complex geometry than a face and also more specularities. Its first five eigenvectors, see Figure 5, are somewhat similar to those obtained for the ball and face, but their contribution to the total variance, see Table 1, is less due to the greater geometric interpolation.

4 Objects With Specular Spikes

We examined two extremely specular objects: a motorcycle helmet and a fire extinguisher. Both objects have almost "mirror-like" surfaces. Most of the images contain diffuse specularities and several of them have very bright specular spikes which are practically mirrored reflections of the light source. See, for example, the left image of Figure 6. It is precisely these specular spikes which will cause trouble for an eigenspace approach because they will change their location and shape dramatically as the lighting direction is changed.

The left image in Figure 6 is a member of the dense set, but not of the sparse set. In the center is the projection onto the first 18 eigenvectors calculated from the dense set. The image is well reproduced, which is not surprising since the image was a member of the training set from which the eigenvectors were calculated, and the first 18 eigenvectors account for 92% of the variance. (See Table 1.) However, when the image is projected onto the first 18 eigenvectors calculated from the sparse set, the results, as shown on the right, are terrible. Not only does the projection fail to represent the specular spike, it contains a faint "ghost" specular spike in the wrong location! Figure 7, which shows the first five eigenvectors calculated from the sparse set, shows the source of the problem. The sparse set does not contain an image with the specular spike in the center of the visor, so its eigenvectors cannot reproduce this spike. However, the sparse set does include an image with a specular spike near the top of the visor, and this particular spike is represented prominently in the fifth eigenvector. So when we attempt to project the image in Figure 6 onto the eigenspace, the reconstruction shows a "ghost" specular spike instead of the real specular spike. The image in Figure 6 and the image in the sparse dataset are similar to each other–the difference in lighting direction between them is only 15 degrees. However, this change in lighting direction has caused a significant "movement" of the specular spike–and the linear combination model is unable to make a geometric interpolation.

Figure 8 shows the eigenvectors calculated from the dense set, with lighting from both sides. The first five eigenvectors do roughly correspond to those for the face, ball, and parrot. But their contribution to the total variance is small, see Table 1, because of the large amount of specularities. Something very interesting happens with the higher eigenvectors here. Rather than representing single specular spikes, these eigenvectors tile the space in a regular way such that spikes at different locations can be represented (one is reminded of a Fourier transform.) This general representation scheme for single specular spikes was observed again in the eigenvectors for the fire extinguisher. (See [10].) We have a nice progression from wide specular lobes in the lower eigenvectors to narrower specular lobes in the higher eigenvectors.

Of course, not all lighting conditions result in a single specular spikes. Figure 9 shows a picture of the helmet taken under ambient lighting conditions. Instead of one specular spike, we now have several. The middle image shows what happens when we project this image onto the first 5 eigenvectors– the specular spikes are not very well reproduced at all.

This implies that representing specularities will require a very large number of eigenimages. Accordingly, since the goal is object recognition and not image representation, it is preferable to use the first few eigenimages to represent the Lambertian and specular lobes, while treating the specular spikes as residuals. If these residuals are sufficiently small then they will have little effect on matching. It might also be possible to obtain a probabilistic model for the specularities themselves in isolation, but we will not attempt this here.

But can we find eigenvectors that describe the Lambertian and specular lobe components but which ignore the spikes? Observe in Figure 8 that the lower eigenvectors represent spatially diffuse lighting effects– only the higher eigenvectors represent specular spikes. As Figure 9 demonstrates, projecting an image with specular spikes onto the first five eigenvectors largely eliminates specular spikes–at least, they no longer dominate the image. Instead they show up in the residuals.

Thus, we can characterize the first five eigenvectors as corresponding to the Lambertian and specular lobe components. The remaining eigenvectors corresponding to the specular spikes, which are best treated as residuals. This characterization is only approximate but, encouragingly, the bigger the training set the more accurate it appears to be – it is far more accurate for the dense model than for the sparse model.

5 Objects With Shadows

Our results for faces suggest that sharp shadows might be a problem for our model. To investigate this, we examine several objects with surface geometries that cast distinct shadows. One of these objects, a function generator–shown in Figure 11–has two interesting characteristics: a bright but quite localized specularity on the big knob, and sharp shadows cast by the other knobs.

Figure 10 shows the first 10 eigenvectors calculated from the sparse set for this object. The first few eigenvectors have similar interpretations to those of the ball and the face–as frontal, top, side, and extreme-side lighting. Because there is only one specularity –on the knob–these eigenvectors can reproduce it. Higher eigenvectors capture more precisely the shadows cast by the smaller knobs.

Figure 11 shows the results of projecting an image from the dense set onto the first 5 and the first 10 eigenvectors calculated from the sparse set. Although the specularity on the knob is well represented in the reconstruction from the projection onto 5 eigenvectors, the small shadows are lost. The shadows start to appear in the reconstruction from the projection onto 10 eigenvectors, though they are not quite correct. Nevertheless, in neither of the projections does the absence of correct small shadows change the general appearance of the object very much. Perceptually, small shadows do not seem to be very salient.

The amount of variance accounted for by the first 10 eigenvectors is shown in Table 1. Note that the first eigenvector accounts for almost 80% of the variance! This is because the object is, on a rough scale, just a flat lambertian surface. The images formed by such a surface under different lighting conditions are all just multiples of a single vector.

Another object with distinct shadows–an infrared detector–is shown in Figure 12. This object has two distinct types of shadow– a large shadow cast by the handle on the top right, and some local shadows cast by the knobs at the bottom. The second image shows a reconstruction from the projection of the original image onto the first five eigenvectors calculated from the dense set. This is too few eigenvectors to reconstruct the shadow cast by the handle. Nevertheless, the residuals, shown in the third image, are not very prominent.

The fourth image shows the reconstruction from projection onto the first 18 eigenvectors calculated from the dense set. The geometry of the shadow cast by the handle is correctly reproduced, but the shadow is not as distinct as it was in the original image. The fifth image shows the results of projection onto 5 eigenvectors from the sparse model. The input image is not a member of the sparse set, so we would expect that these eigenvectors will not correctly reproduce the geometry of the shadow–in fact, without the higher eigenvectors, the shadow can hardly be reproduced at all. The last image in Figure 12 shows the results of projecting onto the first 18 eigenvectors calculated from the sparse set. We do get a sharp shadow–however, careful examination of the shadow shows that its shape is different in this projection. This is exactly the same phenomenon we saw with the specularity on the helmet–projection onto a set of eigenvectors which was not calculated using the image in question can result in spurious specularities and shadows–shadows which are not quite in the right place. Note, however, that the spurious shadow is much less noticeable than the spurious specularity. Note also that, for all the projections in Figure 12, the shadows near the knobs are more or less correctly reproduced. The implication is as follows: even though one gets the same geometric problems with shadows as one gets with specularities, the shadows are much less salient and much less of a problem.

This is borne out in Figure 13, which shows the results of projecting an ambient image of this object onto the eigenvectors. On the left is the original image. In the center is the result of projecting onto 5 eigenvectors from the sparse set. On the right is the result of projecting onto 18 eigenvectors from the dense set. As can be seen, both projections are reasonable. This is because, under ambient lighting conditions, shadows are not very prominent.

We conclude that, as for specular objects, only the lower eigenimages need to be used to construct the model.

6 Residual Analysis

In the previous sections, we have shown that five eigenvectors are consistently sufficient to form a useful representation of the lighting of an object. This claim was based on the following observations: (1) lambertian and specular-lobe reflectance can be well represented in a five-dimensional space, (2) specular spikes can be eliminated by projecting onto this space, and (3) localized shadows are not prominent enough to cause problems. In order to further support this contention, it will be instructive to examine the residuals formed by taking the difference between the original image and its projection onto the five-dimensional space. Our claim, then, is that the five eigenvector model is a good representation because these residuals are (1) small, and (2) correspond to easily interpreted features of the lighting, such as specularities and shadows. This claim was illustrated for faces by Figure 4.

The residuals for the helmet are shown in Figure 9. As can be seen the specularities stands out prominently. Most of the rest of the image is fairly well matched. Despite the specularity, the average residual is quite small. The mean pixel difference between the two images is only 5.8 greyscale levels.

This demonstrates that the low-dimensional model can be useful for the identification and separation of specular spikes, sharp shadows, and "transient" features. In a sense, these phenomena can be thought of as "painted over" the underlying image. A natural extension of this would be to use the low-dimensional model to distinguish occluding objects from the underlying image. Figure 14 shows how this is done. On the left is the original image–the helmet with a large occluding object. The second image is the result of projecting this onto 5 eigenvectors. Note that both the occluder and the specularity have been eliminated in the projection. The third image is the residual, which largely identifies the occluder and the specularity. The right image is the result of thresholding the third image such that pixels with intensity differences of greater than 35 greyscale levels are black. The specular spike and the occluder can be easily isolated.

One might imagine that local shadows could be localized and eliminated in the same way, but, as discussed above, this is not the case. The first three images Figure 12 show the results of projecting an image with shadows onto five eigenvectors. The residuals are slight– the mean value of the difference between the original image and the projection is 4.4 greyscale levels. Though the missing shadow on the top right shows up in the residual, it is not very prominent. We conclude that residuals for shadows are too slight to be of much bother.

The fact that we can isolate specularities by using a low-dimensional eigenspace suggests a further extension of this work. The specularities on the helmet are compact geometric entities, therefore it may be possible to come up with a geometric model for the specularities. This may allow the specularities to be used as a cue for object recognition, since the specularities on many objects are confined to distinct regions of the object (recall the specularity on the knob of the voltmeter in Figure 11). In order to better separate specularities from the rest of the image, a method of robust statistics would be useful. In this case, the

specularities would be eliminated as outliers, while the rest of the image is projected onto the eigenspace. We experimented with such a system, see Figure 15. Although we found it straightforward to extract specularities, using the EM algorithm, we did not find any advantage in distinguishing between specularity and non-specularity during projection. The results obtained using the robust model were practically indistinguishable from those using the standard model. It is, however, probably desirable to represent the extracted specularities using a geometric model, but this is beyond the scope of this paper.

7 Conclusion

We have shown that, for a range of objects, only a small number of eigenimages are needed to approximately represent the intensity changes caused by variations in the lighting conditions. The precise number of eigenimages required depends on the type of object. For the degenerate case of a flat, lambertian surface, one eigenvector suffices, while for very simple curved objects without specularities three will suffice. When the object has specular lobes, additional eigenvectors are needed. In the extreme case when the object is very highly specular, no low dimensional model exists that can reconstruct all the highlights accurately. However, provided the sampling of lighting directions in the training set is dense, a low dimensional model can be built that does capture the lambertian and specular lobes. Specular spikes, sharp shadows and occluders can be treated as residuals in this model.

These results are demonstrated by visual inspection of the reconstructions, shown in the figures, and by plots of the goodness of fit and by analysis of the residuals. These residuals are typically small and can be easily identified as specular spikes, sharp shadows or occluders.

We summarize the results with the heuristic that, for a wide range of images, 5 ± 2 eigenvectors suffice to build low dimensional spaces that accurately capture the intensity variations caused by different lighting conditions. Moreover, the number of training images required is relatively small (lighting directions can be sampled every 30 degrees) for objects without much specularity or self-shadowing. Though such a low dimensional model cannot be used to reconstruct intensity features such as specular spikes and small, sharp cast shadows that vary geometrically as the lighting direction changes, the model can be used to eliminate and/or localize such features. Thus, our lighting model eliminates specular spikes and ignores sharp shadows, while representing the gross shading correctly. For purposes of object recognition, we can therefore use a low dimensional model and treat specular spikes, sharp shadows, and occluders as residuals (See [4]).

Acknowledgments

We would like to thank David Mumford for helpful discussions. This research was supported in part by the Brown/Harvard/MIT Center for Intelligent Control Systems with U.S. Army Research Office grant number DAAL03-86-K-0171. The authors would also like to thank ARPA for Air Force contract F49620-92-J-0466 and the NSF for grant IRI-9317670.

References

[1] P. W. Hallinan, "A low-dimensional representation of human faces for arbitrary lighting conditions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 995–999, 1994.

[2] Y. Adini, Y. Moses, and S. Ullman, "Face recognition: the problem of compensating for changes in illumination direction," Tech. Rep., Department of Applied Mathematics and Computer Science, Weizman Institute of Science, 1993.

[3] P. W. Hallinan, "A low-dimensional representation of human faces for arbitrary lighting conditions," Tech. Rep. 93-6, Robotics Laboratory, Harvard University, 1994.

[4] P. W. Hallinan, "A deformable template for recognizing human faces under arbitrary lighting conditions," Tech. Rep., Robotics Laboratory, Harvard University, 1995. In preparation.

[5] Y. Moses. personal communication.

[6] A. Shashua, *Geometry and Photometry in 3D Visual Recognition*. PhD thesis, MIT Dept of Brain and Cognitive Science, August 1992.

[7] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: physical and geometric perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 611–634, 1991.

[8] D. Forsyth and A. Zisserman, "Reflections on shading," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, 1991.

[9] A. Pentland, T. Starner, N. Etcoff, A. Masoiu, O. Oliyide, and M. Turk, "Experiments with eigenfaces," in *Looking at People Workshop, Proceedings of IJCAI'93, Chamberry, France*, 1993.

[10] R. Epstein, P. W. Hallinan, and A. L. Yuille, "5±2 eigenimages suffice: An empirical investigation of low-dimensional lighting models," Technical Report 94-11, Robotics Laboratory, Harvard University, 1994.

[11] D. Mumford, "Parameterizing exemplars of categories," *Journal of Cognitive Neuroscience*, vol. 3, p. 88, 1991.

Eigen-vector	ball	parrot	phone	face	helmet	helmet	helmet	fire ext	function generator	infrared detector
	sparse (right)	dense (both)	dense (right)	sparse (right)	sparse (right)	dense (right)	dense (both)	dense (right)	dense (right)	dense (right)
#1	0.482	0.428	0.679	0.537	0.320	0.340	0.388	0.536	0.806	0.624
#2	0.844	0.697	0.832	0.752	0.569	0.540	0.474	0.687	0.879	0.805
#3	0.944	0.763	0.882	0.902	0.651	0.628	0.581	0.765	0.922	0.885
#4	0.965	0.815	0.920	0.921	0.728	0.746	0.655	0.816	0.936	0.915
#5	0.979	0.847	0.941	0.935	0.798	0.772	0.722	0.852	0.948	0.927
#6	0.989	0.872	0.952	0.945	0.845	0.794	0.750	0.882	0.956	0.939
#7	0.991	0.885	0.963	0.953	0.881	0.816	0.795	0.901	0.962	0.948
#8	0.993	0.897	0.968	0.958	0.905	0.833	0.811	0.913	0.968	0.954
#9	0.995	0.907	0.972	0.963	0.924	0.848	0.824	0.925	0.972	0.960
#10	0.996	0.917	0.975	0.966	0.943	0.861	0.837	0.933	0.975	0.965

Table 1: The variance (cumulative) accounted for by each eigenvector for several different objects, both for sparse and dense training sets.

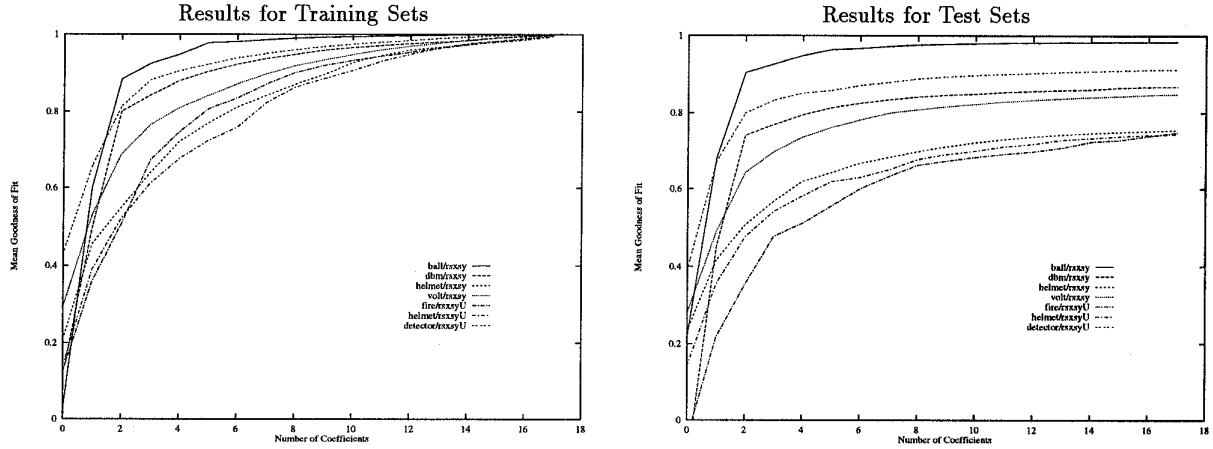

Figure 1: Plots of the mean goodness of fit vs the number of eigenvectors for training and test sets of several different objects. All models were constructed using a sparse data set. The letter 'U' appended to an object's name indicates the background was unmasked.

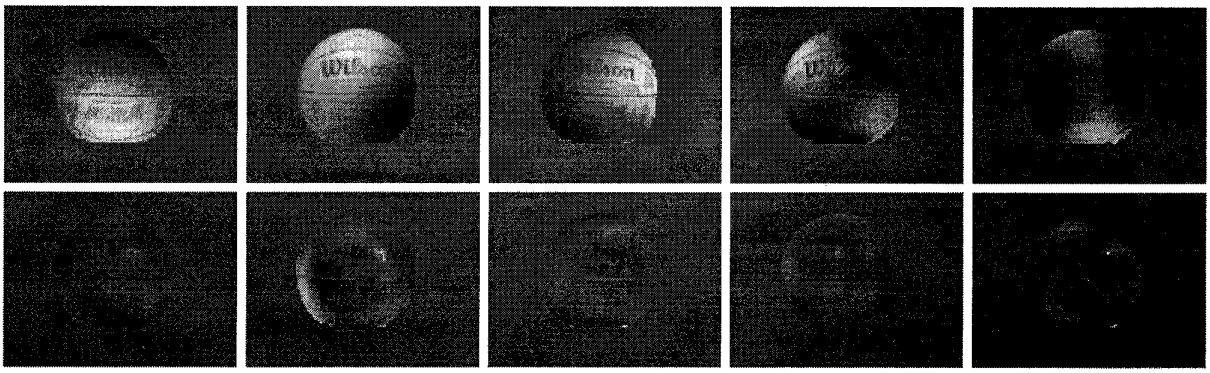

Figure 2: Eigenvectors calculated from the sparse set for the ball, arranged from left to right, top to bottom, in order of importance as determined by their eigenvectors.

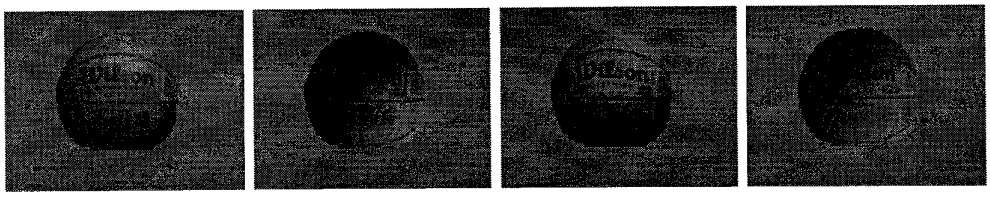

Figure 3: Left two images are data images from the dense set for the ball. These images were not used to calculate the eigenvectors. Right two images are reconstructions of the left two images from their projections onto the first 3 eigenvectors calculated from the sparse set.

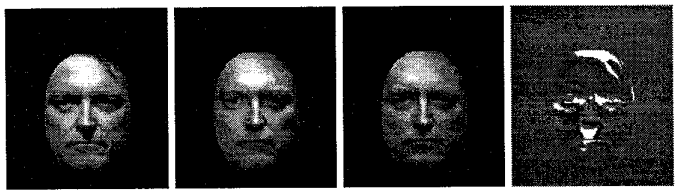

Figure 4: On the left is an image of a face from the dense set. The second image is the reconstruction from the projection onto the first 3 eigenvectors calculated from the sparse set. The third image is the reconstruction from its projection onto the first five eigenvectors calculated from the sparse set. The rightmost image shows the difference between the first and third images, where differences of greater than 30 greyscale levels have been highlighted in black (if the original image is darker than the reconstruction from the projection), or white (if the original image is brighter than the reconstruction from the projection).

Figure 5: First five eigenvectors calculated from the dense set for an artificial stuffed parrot.

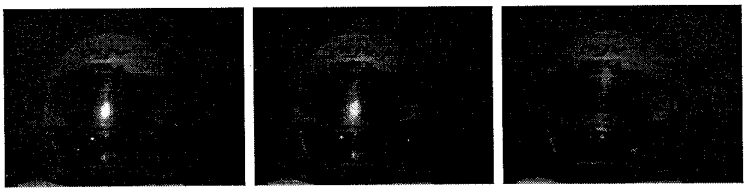

Figure 6: On the left is an image from the dense set. In the center is the reconstruction from the projection onto the first eighteen eigenvectors calculated from the dense set. The image is well reproduced. On the right is the reconstruction from the projection onto the first 18 eigenvectors calculated from the sparse set, of which the image is not a member. The specular spike cannot be reproduced in the correct location.

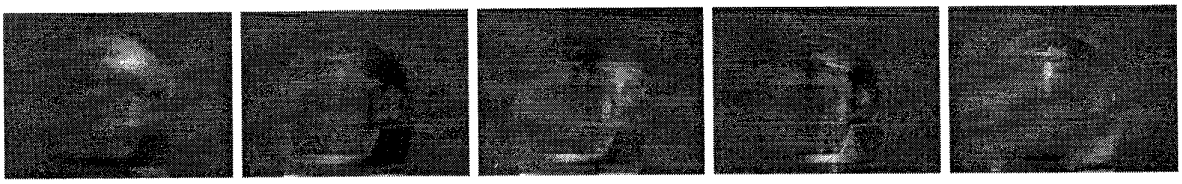

Figure 7: First five eigenvectors calculated for the helmet from the sparse set of images. Lighting was varied over the right side only

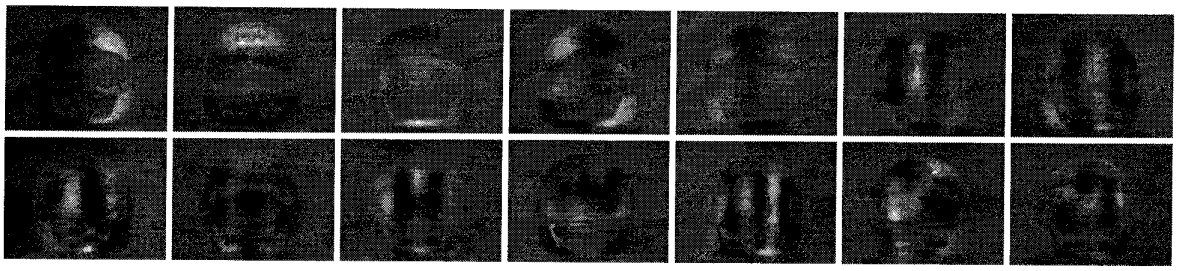

Figure 8: Eigenvectors calculated for the helmet from the dense set; lighting from both sides.

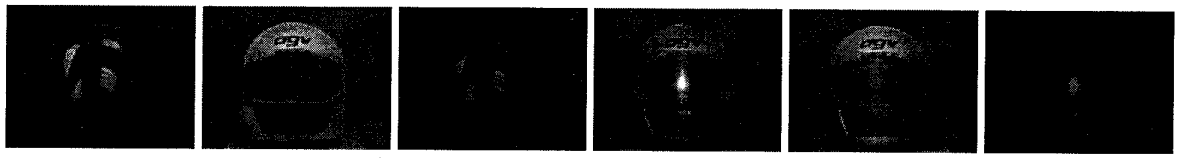

Figure 9: First image is helmet under ambient lighting conditions. Second image is reconstruction from projection onto first 5 eigenvectors. Note the specularities have been eliminated. Third image is the difference between original image and reconstruction. (Differences less than zero have been set to zero.) The specularities are easily isolated. Fourth, fifth and sixth image show the same effect for an image from the dense data set.

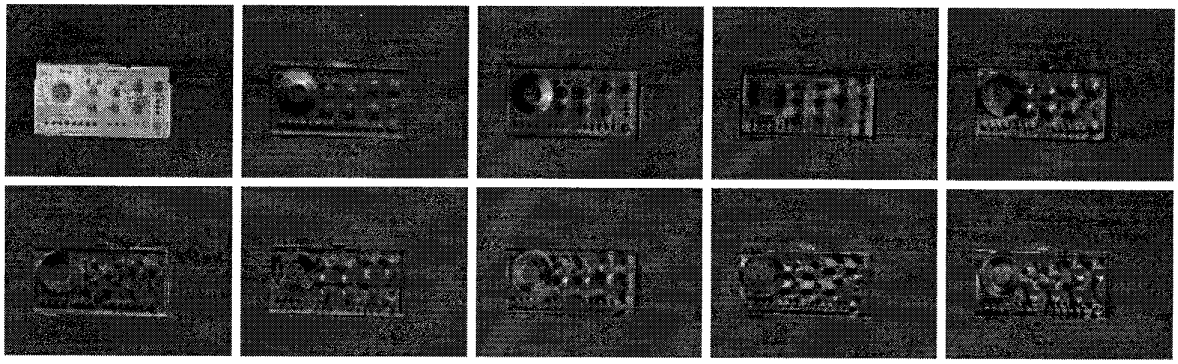

Figure 10: Eigenvectors calculated from the sparse set.

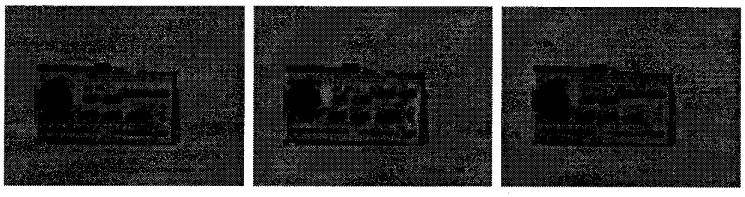

Figure 11: Left is a data image from the dense set. Center is its reconstruction from projection onto first five eigenvectors calculated from the sparse set. Right is reconstruction from projection onto first ten eigenvectors calculated from the sparse set.

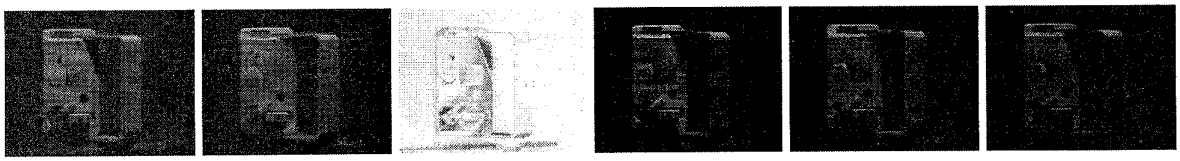

Figure 12: On the left is the original image from the dense set. Second image is reconstruction from projection onto the first 5 eigenvectors from the dense set. Third image shows the difference between the first two images. The residuals are very slight. Fourth image is reconstruction from projection onto the first 18 eigenvectors from the dense set. Fifth image is reconstruction from projection onto the first 5 eigenvectors from the sparse set. Sixth image is reconstruction from projection onto the first 18 eigenvectors from the sparse set. See text for commentary.

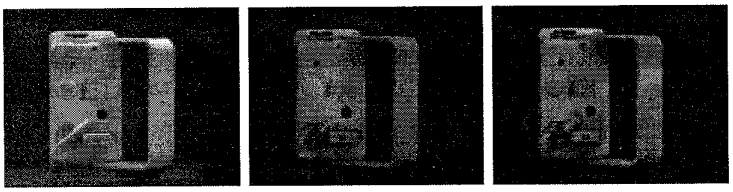

Figure 13: On the left is the IR detector under ambient lighting conditions. Center is reconstruction from projection onto 5 eigenvectors calculated from the sparse set. Right is projection onto 18 eigenvectors calculated from the dense set.

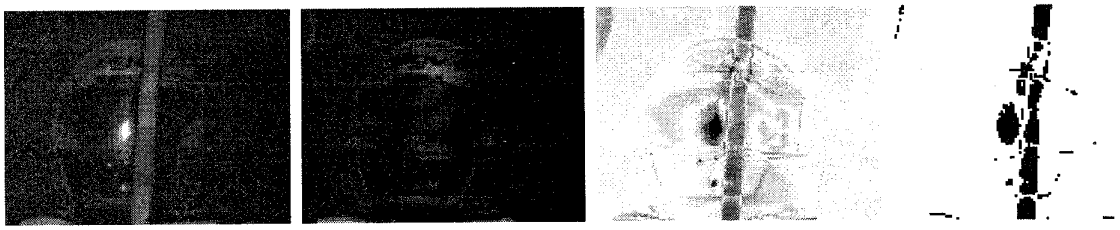

Figure 14: The left image is the helmet with an occluding object. The second image is the result of projecting onto five eigenvectors. Both the occluder and the specularity have been eliminated. The third image shows the residual– the absolute value of the difference between the first and second images. The rightmost image shows the result of thresholding the residual. Black pixels have a difference of more than 35 greyscale levels. The specular spike and the occluder are well isolated.

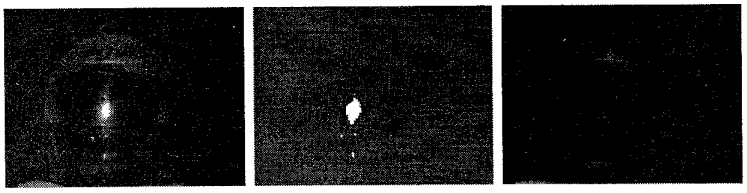

Figure 15: On the left is the original image. The center shows the outliers extracted by a robust statistics process; these pixels were not used in the projection onto the eigenspace. Right is the result of the projection. The specularities are easily extracted; however, the result of the projection is no better than in Figure 9.