ADDING SCALE TO CORNERS

PAUL L. ROSIN
School of Computing Science, Curtin University of Technology
Perth, 6001, Western Australia

ABSTRACT

A failing of standard grey-level corner detectors is that they do not extract most of the attributes of a corner apart from its strength. This paper concentrates on determining one of these additional attributes, namely the range of scales that the corner exists over. It is shown that a multi-scale analysis is beneficial in detecting certain classes of corners likely to arise in scenes (e.g. noisy, textured, rounded, blurred).

1. Introduction

A common paradigm in computer vision is the extraction of simple low-level features from the image, followed by further processing based on these features. However, this paradigm accounts in part for the failure of contemporary computer vision systems to operate in a general purpose and robust manner. The simple features provide an impoverished description, making the analysis of the scene based on this inadequate data - which is usually incomplete, imprecise, and ambiguous - extremely difficult. To remedy this Marr proposed describing the image by the primal sketch [8] which contains detailed information about the intensity changes in the image. As Waltz found when interpreting line drawings of block world scenes, additional information (e.g. extra line labels) provided additional constraints allowing his system to converge more rapidly on a less ambiguous solution [15]. Likewise, a richer description of image features would improve computer vision systems. We have previously argued that image features are commonly under-utilised in computer vision [11]. Much of the information present in a feature is not explicitly extracted from the image, while part of the available information is discarded without being used.

This paper describes grey-level corners, which are a popular low-level image feature, and are used for model matching, and for matching between stereo pairs or motion sequences of images to derive the 3D structure of the scene. In addition, corner matching is insensitive to changes in scale, orientation, and partial object occlusion [16]. Most corner detectors assume an idealised corner that is sharply pointed and has straight, steep edges. However, corners rarely appear like this in the real world. Due to manufacturing limitations, wear and tear, streamlining, aesthetics, and so forth, corners are more typically rounded, blunted, blurred, textured, etc. In fact, the only sharp corners in nature are likely to arise when two objects cross or interpenetrate (i.e. transversality [4]). This variety in appearance does not appear to be a problem to the human visual system since at an appropriate scale the corners are usually well defined. This implies that corner detection at a fixed scale is likely to be problematic, and instead should be performed over multiple scales. No single operator can be optimal over all scales. Moreover, corners exist not just at a single scale, but rather over a finite range of scales. Koenderink [6] has stated that objects generally have an inner and outer scale. Beyond the outer scale the object is too small to be adequately distinguished, while the object does not exist as a whole beyond the inner scale.

We asserted that much of the available useful information present in features is not acquired by the standard feature extraction operators in use today. The majority of corner detectors return just a single value measuring the strength of the corner (e.g. [1, 2, 3, 5, 9]). However, there are many more descriptive attributes present in corners. Besides scale, it would be desirable to extract a corner's subtended angle, orientation, edge shape, edge magnitude, edge profile, sharpness, colour, and junction type.

This paper concentrates on the extraction of the scale parameter of corners. Details of techniques for determining some of the other attributes of corners are given in [12]. A multi-scale approach to corner detection has a number of benefits. First, the range of scales that a corner exists over may be useful when matching the corner. Second, it enables more robust feature detection. This is because corners that exist over a large range of scales are likely to be more significant than corners that only exist at one or two levels in the pyramid. In particular, noise and artifacts in the image will not produce a consistent set of corners over scale. Measuring a corner's significance by its lifetime over scale is akin to Witkin's [17] scale-space analysis. Third, by analysing the properties of a corner over scale it may be possible to determine new attributes of the corner. For instance, the detected position of cusp shaped corners will show a systematic shift over scale, and their subtended angle and orientation may also change. This approach is similar to Sjöberg & Bergholm [14] who classified edges as either object or shadow types by their behaviour over scale.

2. Detecting Corners at Multiple Scales

There are two principle approaches to performing the multi-scale analysis. Either the nature of the corner detector can remain unchanged while the size of the image is reduced, or the image size can remain constant while the corner detector operates over larger areas by increasing its mask sizes, etc. We use the first approach since it has several advantages. First, it only requires the standard implementation of a single scale corner detector. This also allows the multi-scale corner analysis to be easily applied to a variety of corner detectors with little effort. Generalising each individual corner detector to work over areas of varying sizes would prove more complicated and time consuming. A second advantage to reducing the image size is that the processing at each scale will run more quickly. The overhead of generating the smaller images is outweighed by the increased speed in analysing smaller images as compared to analysing the full image with large masks.

The most popular technique for analysing images at different scales is to generate a grey-level pyramid [10]. At each level in the pyramid pixels are calculated from either 2×2 non-overlapping windows or 4×4 overlapping windows in the lower level image. Usually pixels are assigned the average of the window, but other functions such as the median have been used. Despite their simplicity, pyramids have the drawback that images at successive levels always reduce in size by a factor of two. However, scale-space is a continuous domain, and as Koenderink has pointed out, it is undersampled using pyramids [6]. When tracking corners over multiple scales coarsely quantising the scale dimension makes the problem of linking corresponding corners in adjacent scales more difficult. It is difficult to generalise the windowing mechanism of pyramids to produce arbitrary reductions in image size. Therefore we generate the image pyramid using bilinear interpolation between each level. Pyramids can now be generated with any amount of tapering by simply changing a single parameter to the interpolation function. This brings the pyramid much closer to the ideal of continuous scale-space. Allowing the scale parameter to be varied by any amount also enables the sampling of scale-space to be varied adaptively. This can be used to resolve ambiguous feature matches across scale by sampling scale-space more finely at these points (e.g. [7]). However, this has not been implemented in our work.

The correspondence between scales is given by: $x_i = x_j r^{i-j} + \frac{r^{i-j}}{2} - \frac{1}{2}$, where x_i is the position at level i corresponding to the position x_j at level j in the pyramid, and r is the ratio in size between adjacent levels in the pyramid. Of course, there is a trade-off between ease of feature tracking and efficiency. All the examples presented here use a pyramid with levels decreasing in size by $\frac{5}{8}$ which was generally adequate to successfully track corners across scale.

The procedure of tracking corners through scale proceeds top-down. Corners are located at their highest scale and mapped onto the adjacent lower level. At this level a 3×3 window about the expected corresponding location is searched for a corner. If the corner is found its position is used to map onto the next lower level image, and the process is repeated. When plotting tracked corners the finest detected scale is used to determine the corner's position since it is likely to be the most accurate.

3. Detecting Different Corner Types

The following figures show synthetic images containing a variety of corner types with the results of applying the multi-scale corner. Each detected corner is marked by two circles, showing the inner and outer scales. Due to limitations of space the results of applying a standard corner detector (the Kitchen/Rosenfeld detector [5]) at a single scale are not included, but can be found in [12]. The multi-scale analysis applied the same corner detector with the same cornerity threshold at multiple scales and linked corresponding corners at adjacent scales as described above. Only corners that were successfully tracked over a certain number of scales (different ranges are used for the different images) are shown. In describing the results the single scale corner detector will be referred to as SSD, and the multi-scale corner detector as MSD.

A simple corner is shown in figure 1, and as expected the SSD has no difficulty in locating the corner providing similar results to the MSD. However, if Gaussian noise is added to the image (figure 2) many spurious corners are detected at the finest scale by the SSD. At coarser scales the noise is eliminated and only a single corner remains. By eliminating corners with small lifetimes the MSD succeeds in identifying the most significant corner. If texture is added not over all the image but just across the boundaries a wiggly corner is generated as shown in figure 3a. The SSD correctly locates the corners along the boundary, as does the MSD (figure 3b). However in addition, the MSD provides the information that the central corner is more significant than the others since it exists over a larger range of scales than the wiggles.

When the original corner is rounded or smoothed (figures 4 & 5) then the SSD fails to find the corner. Within its small window of interest the boundary is locally straight or has a low grey-level contrast. Although the majority of corners will appear rounded and smoothed in images very few corner detectors can cope adequately with these effects. An exception is Davies [2] who explicitly set out to detect the *crumbly* corners of biscuits. The MSD successfully locates both types of corners at a coarser scale. With the rounded corner example the corners at the finer scales are not positioned accurately. This leads to the final tracked corner being misplaced since fine scales have been assumed to be more accurate than coarse ones. With the blurred corner example the determined range of scales does not extend to the finest scale.

In conclusion it can be seen that applying the corner detector at a single scale to certain classes of corners common in images results in many false positives and negatives whereas the multi-scale analysis is able to correctly identify the significant corners.

4. Examples of Multi-Scale Corner Detection

The following figures show an example of applying the multi-scale corner detector to a real image. Naturally the results are dependent on the performance of the underlying single scale corner detector. For

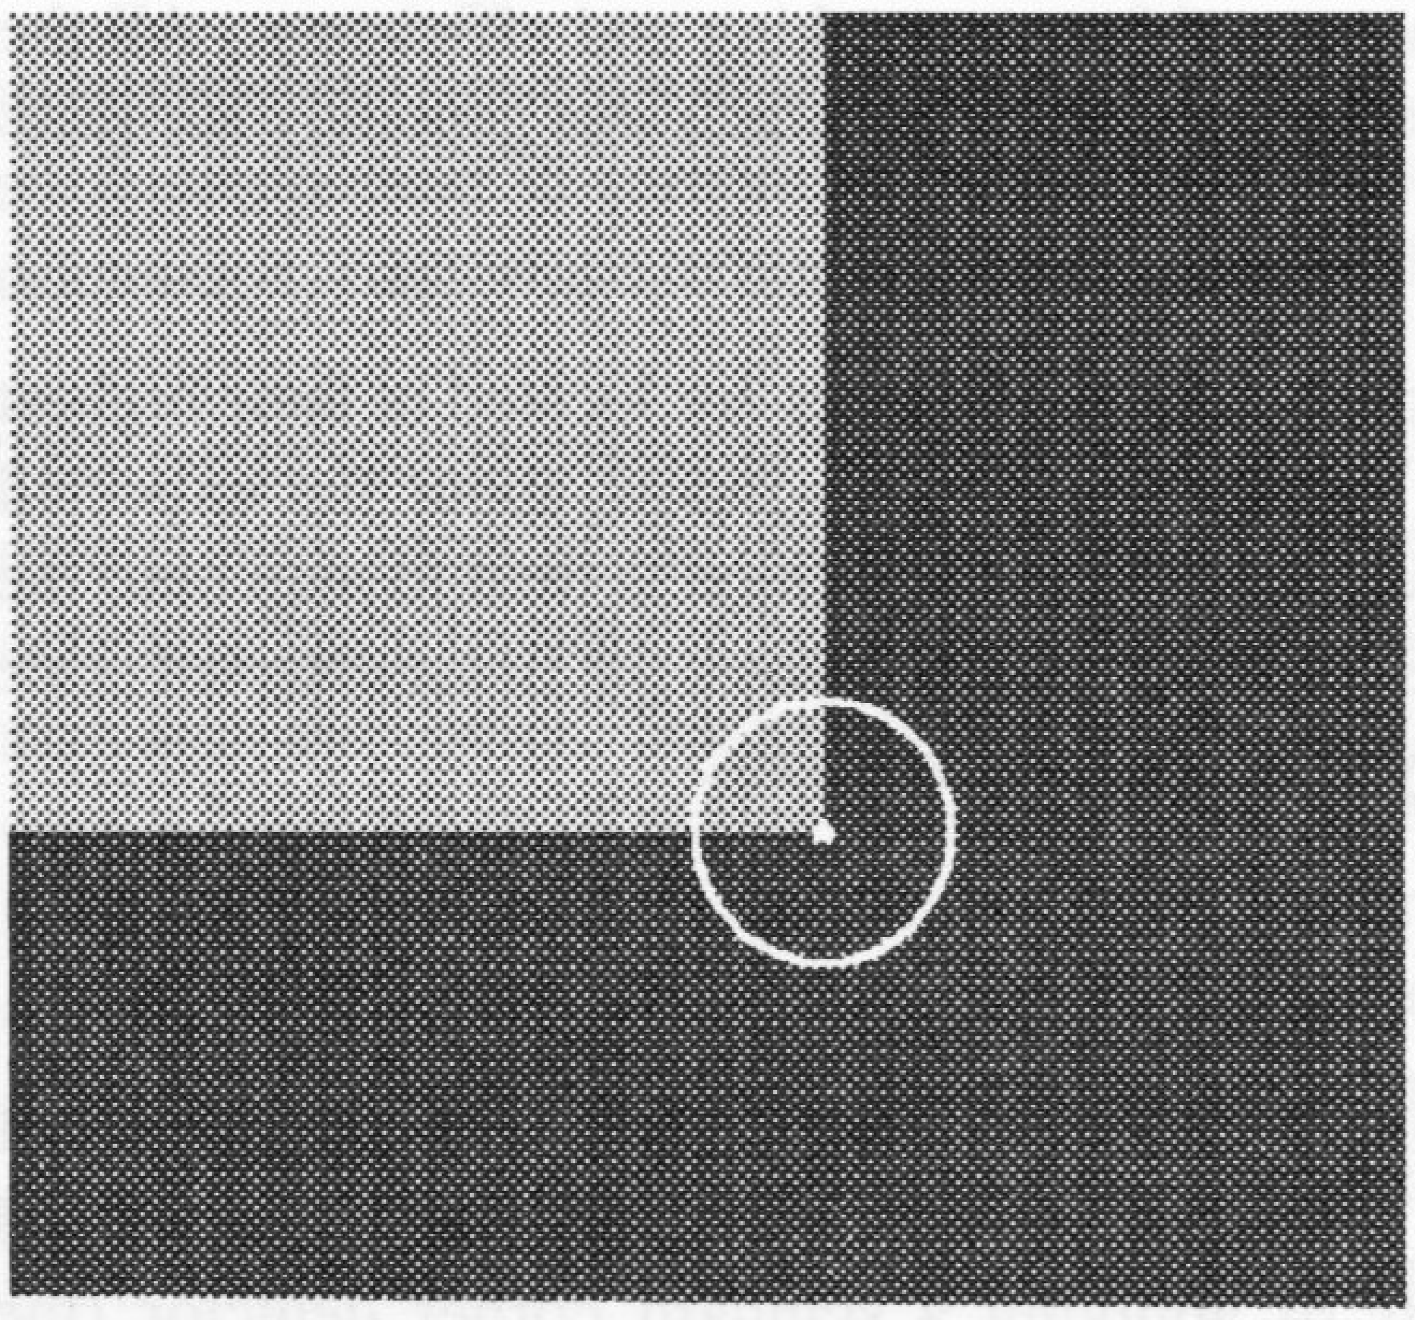

figure 1 - simple corner

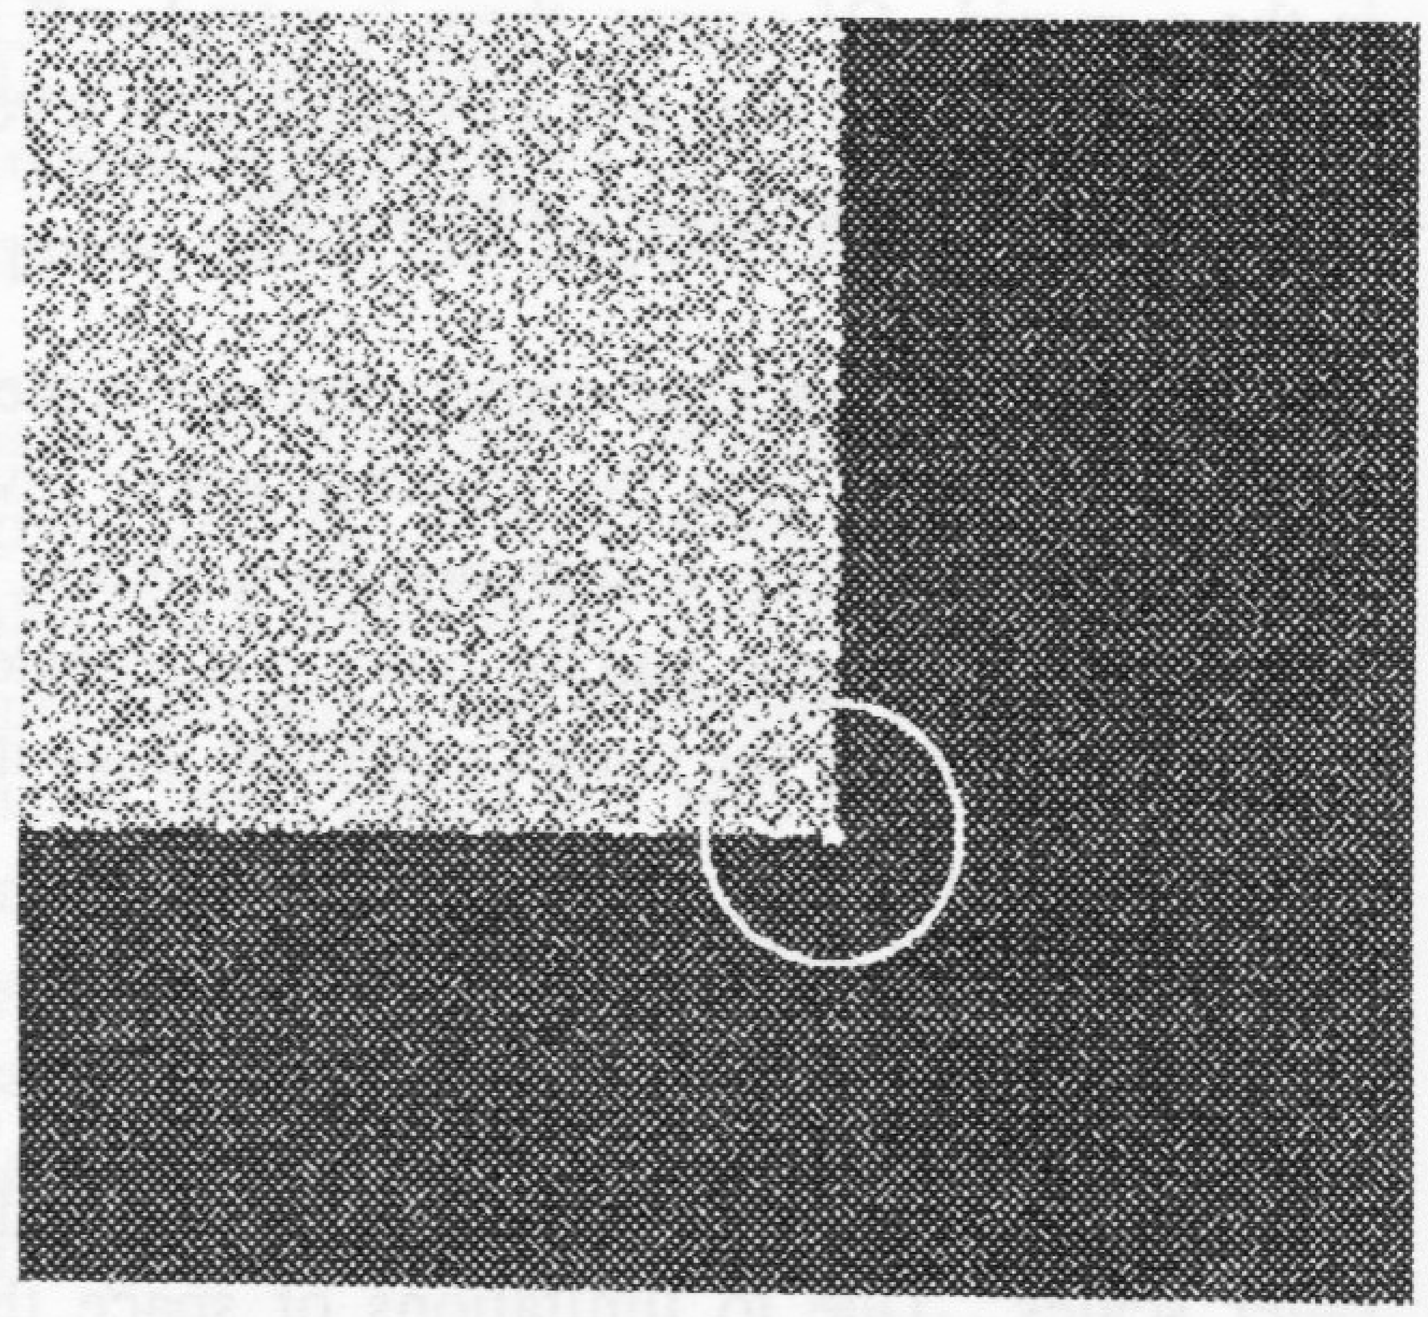

figure 2 - corner with noise added to image

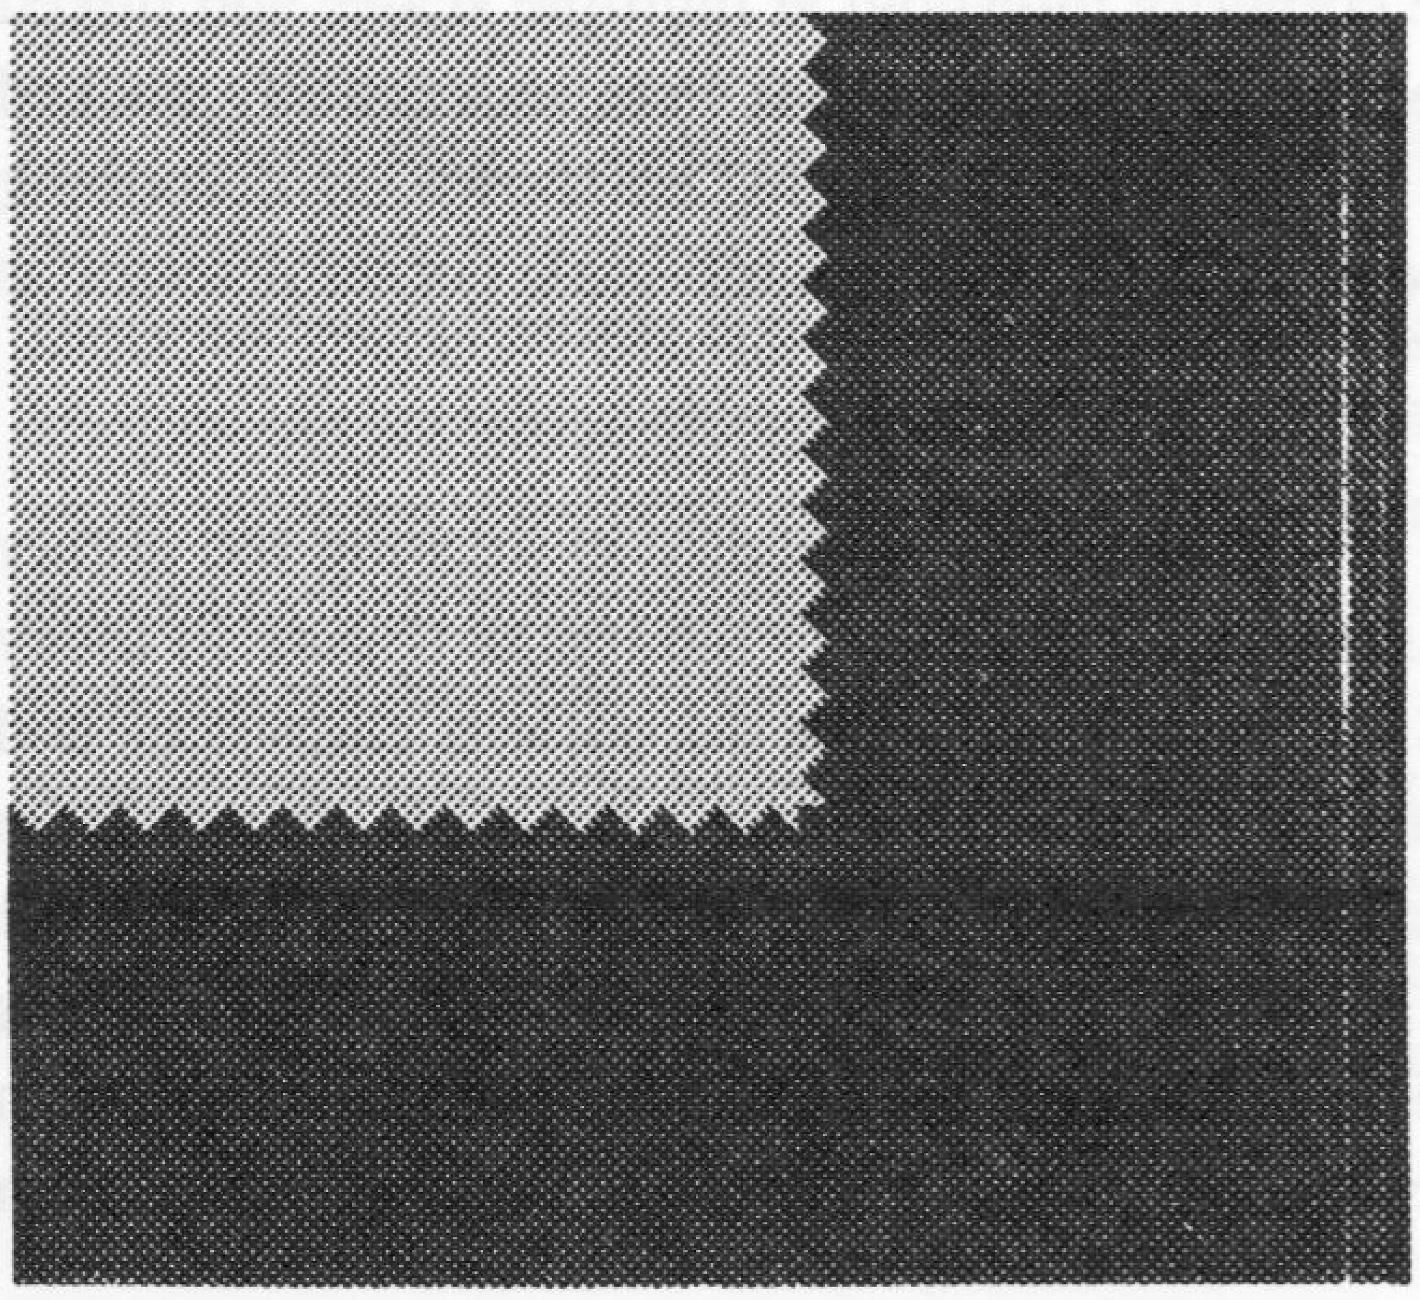

figure 3a - original wiggly boundary image

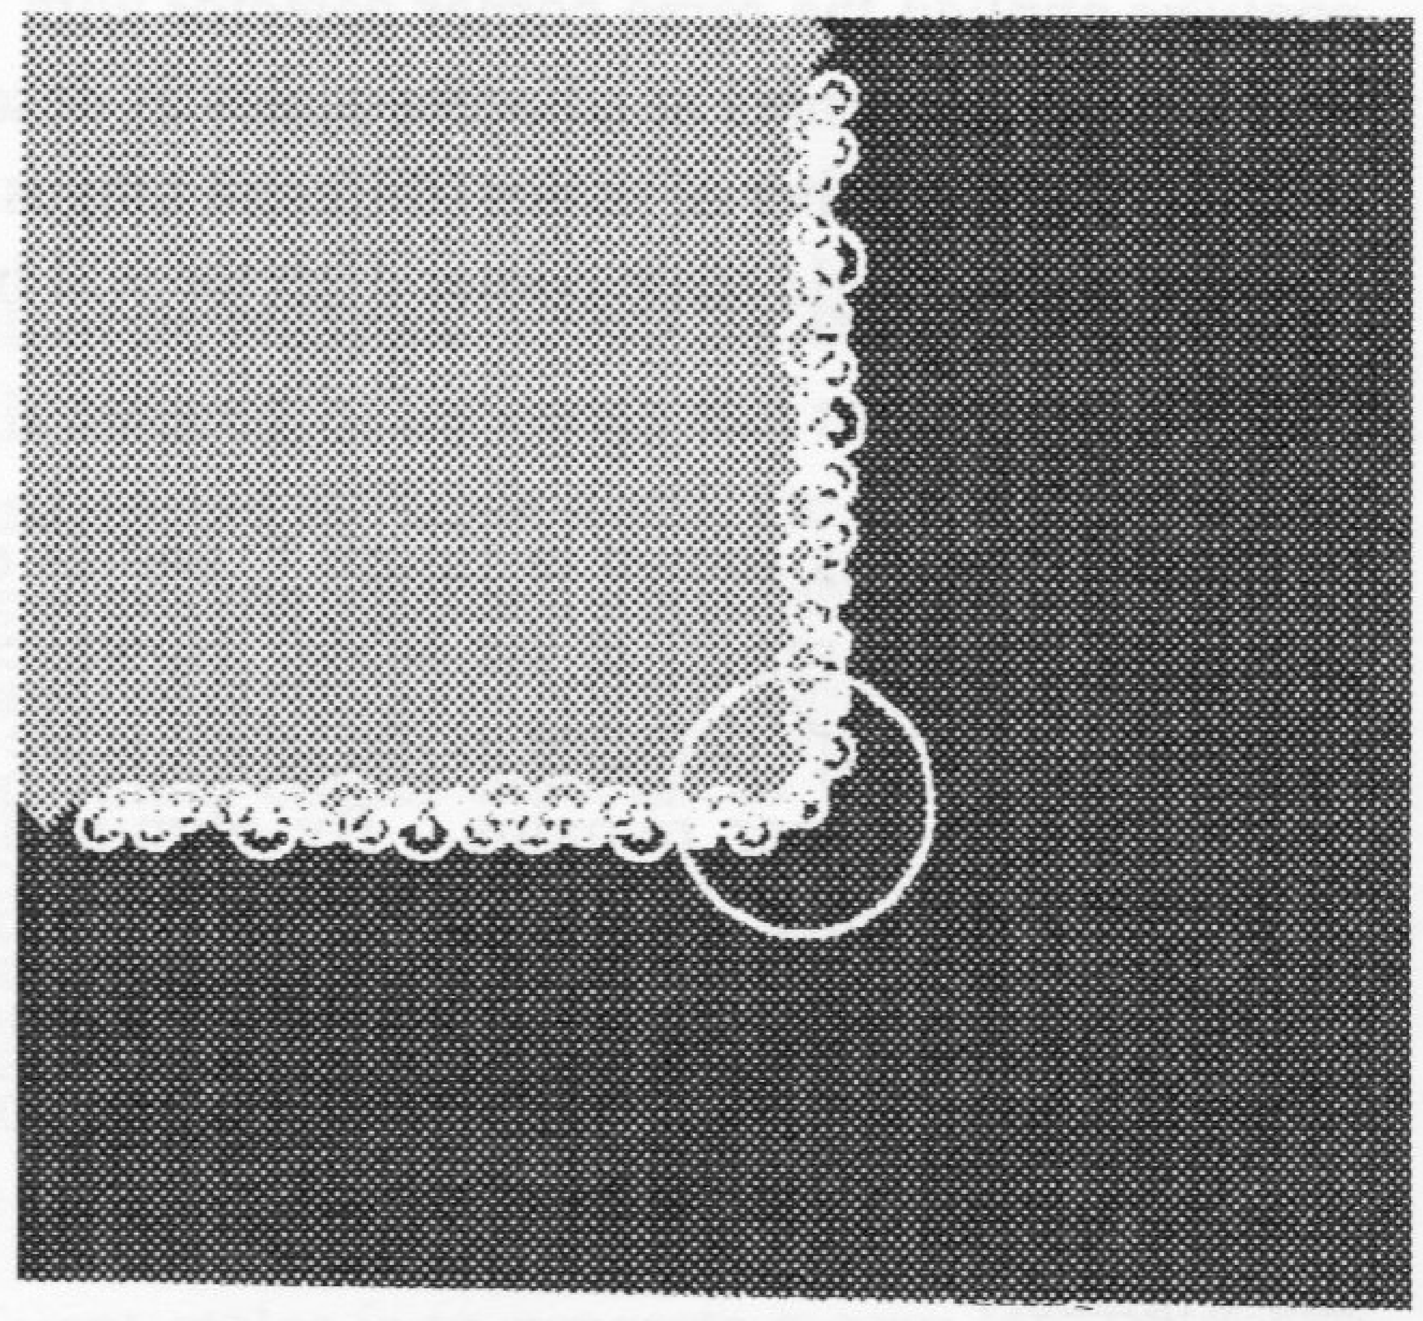

figure 3b - corner detected wiggly boundary

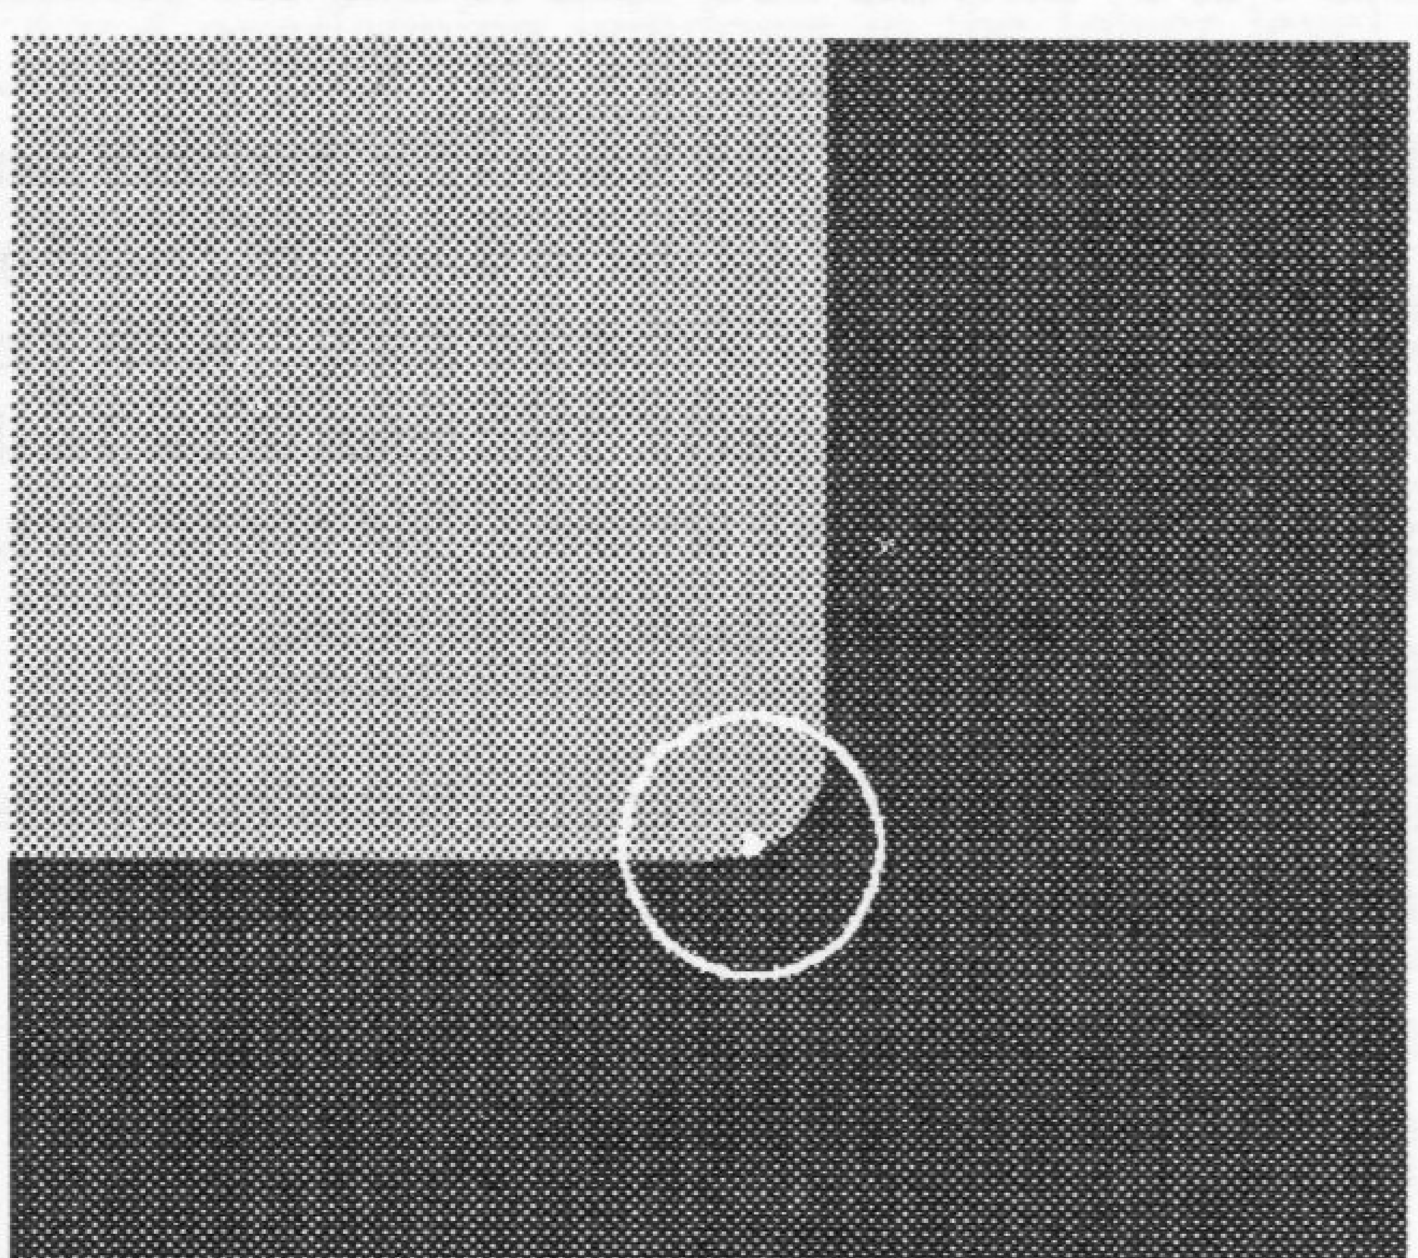

figure 4 - rounded corner

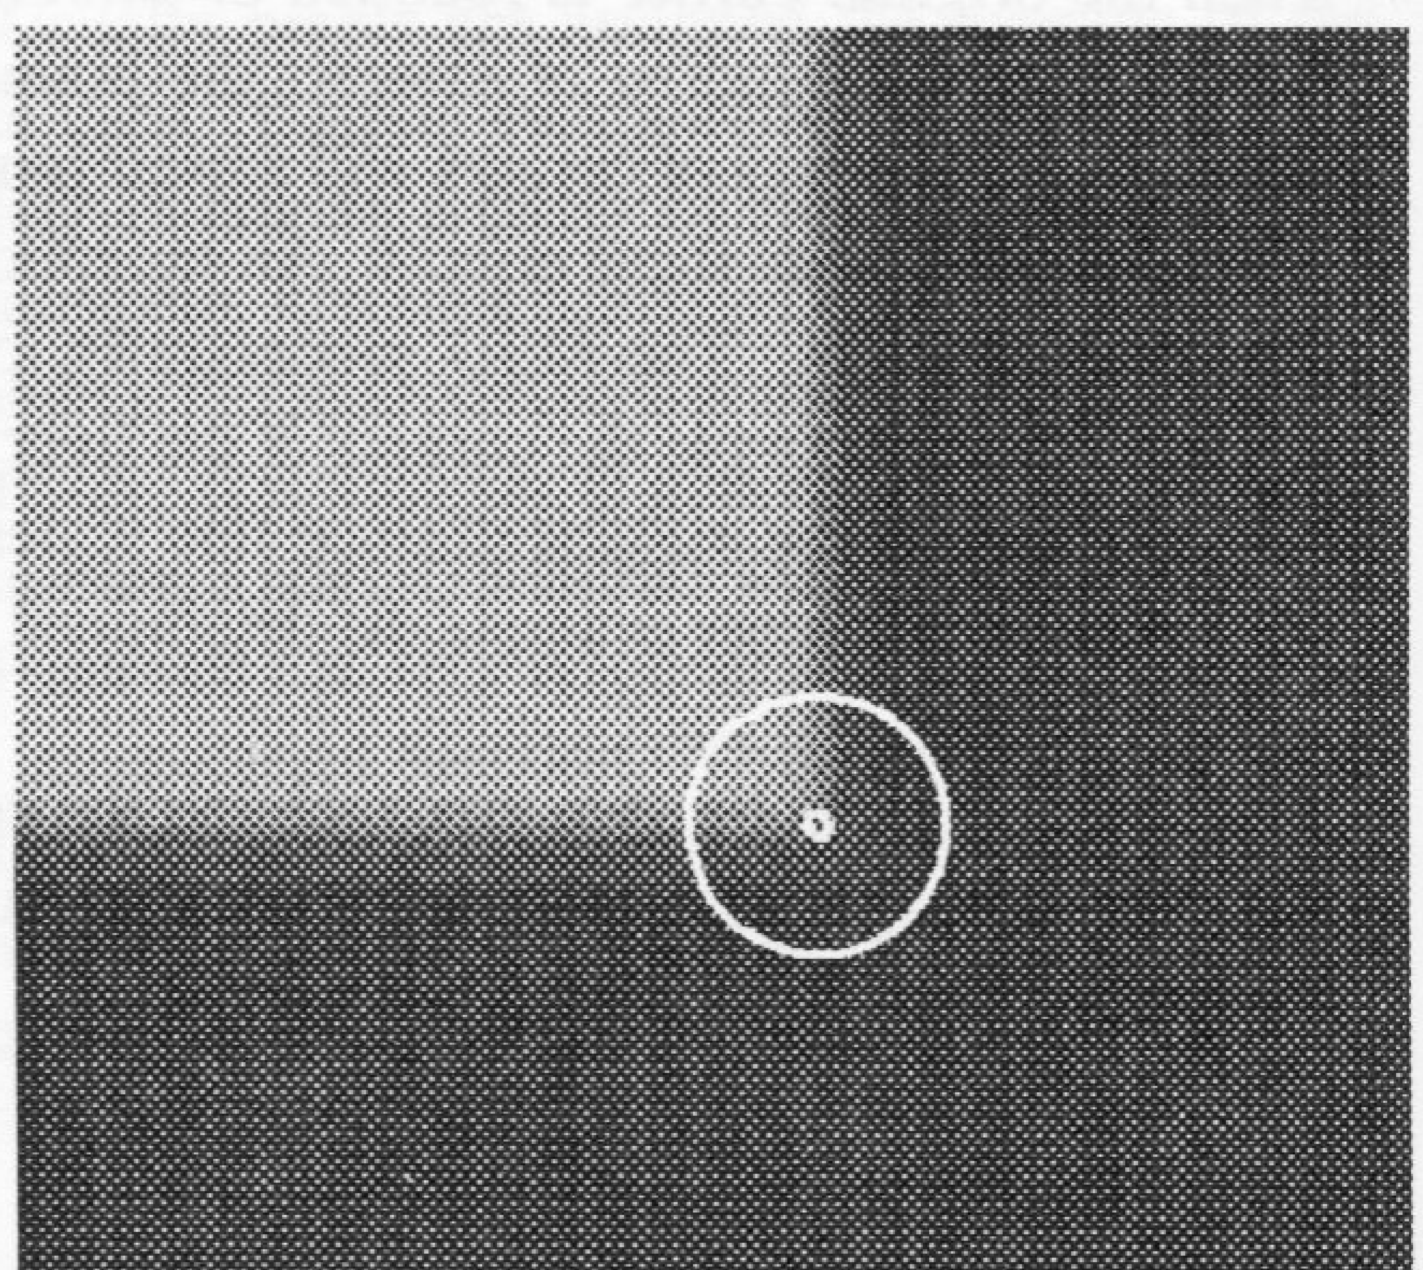

figure 5 - smoothed corner

comparison the results of using several different corner detectors are presented. The cornerity threshold has been set for each detector so that they all successfully track the same number of corners across scales. Only corners which exist over at least four adjacent scales are included. The results of applying the Kitchen/Rosenfeld detector [5] are shown in figure 6, the median filter in figure 7, the Moravec detector [9] in figure 8, and the Plessey detector [3] in figure 9.

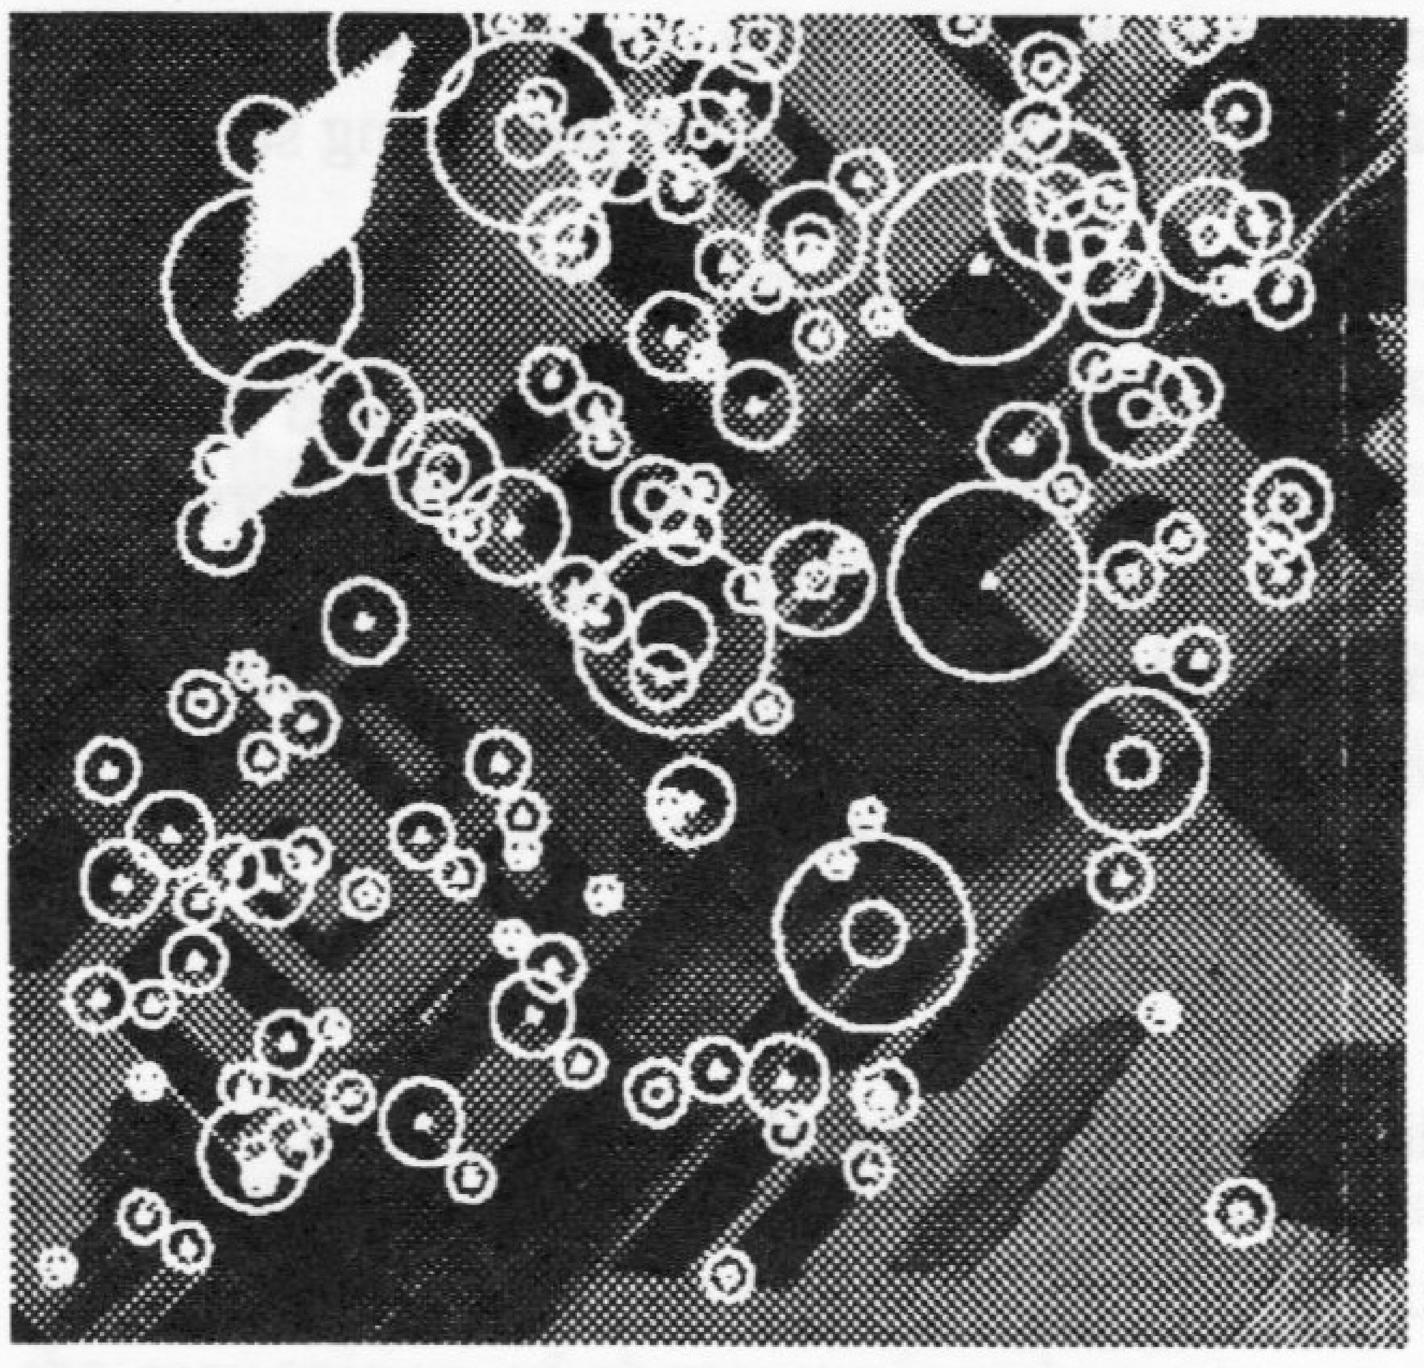

figure 6 - Multi-scale Kitchen/Rosenfeld detector

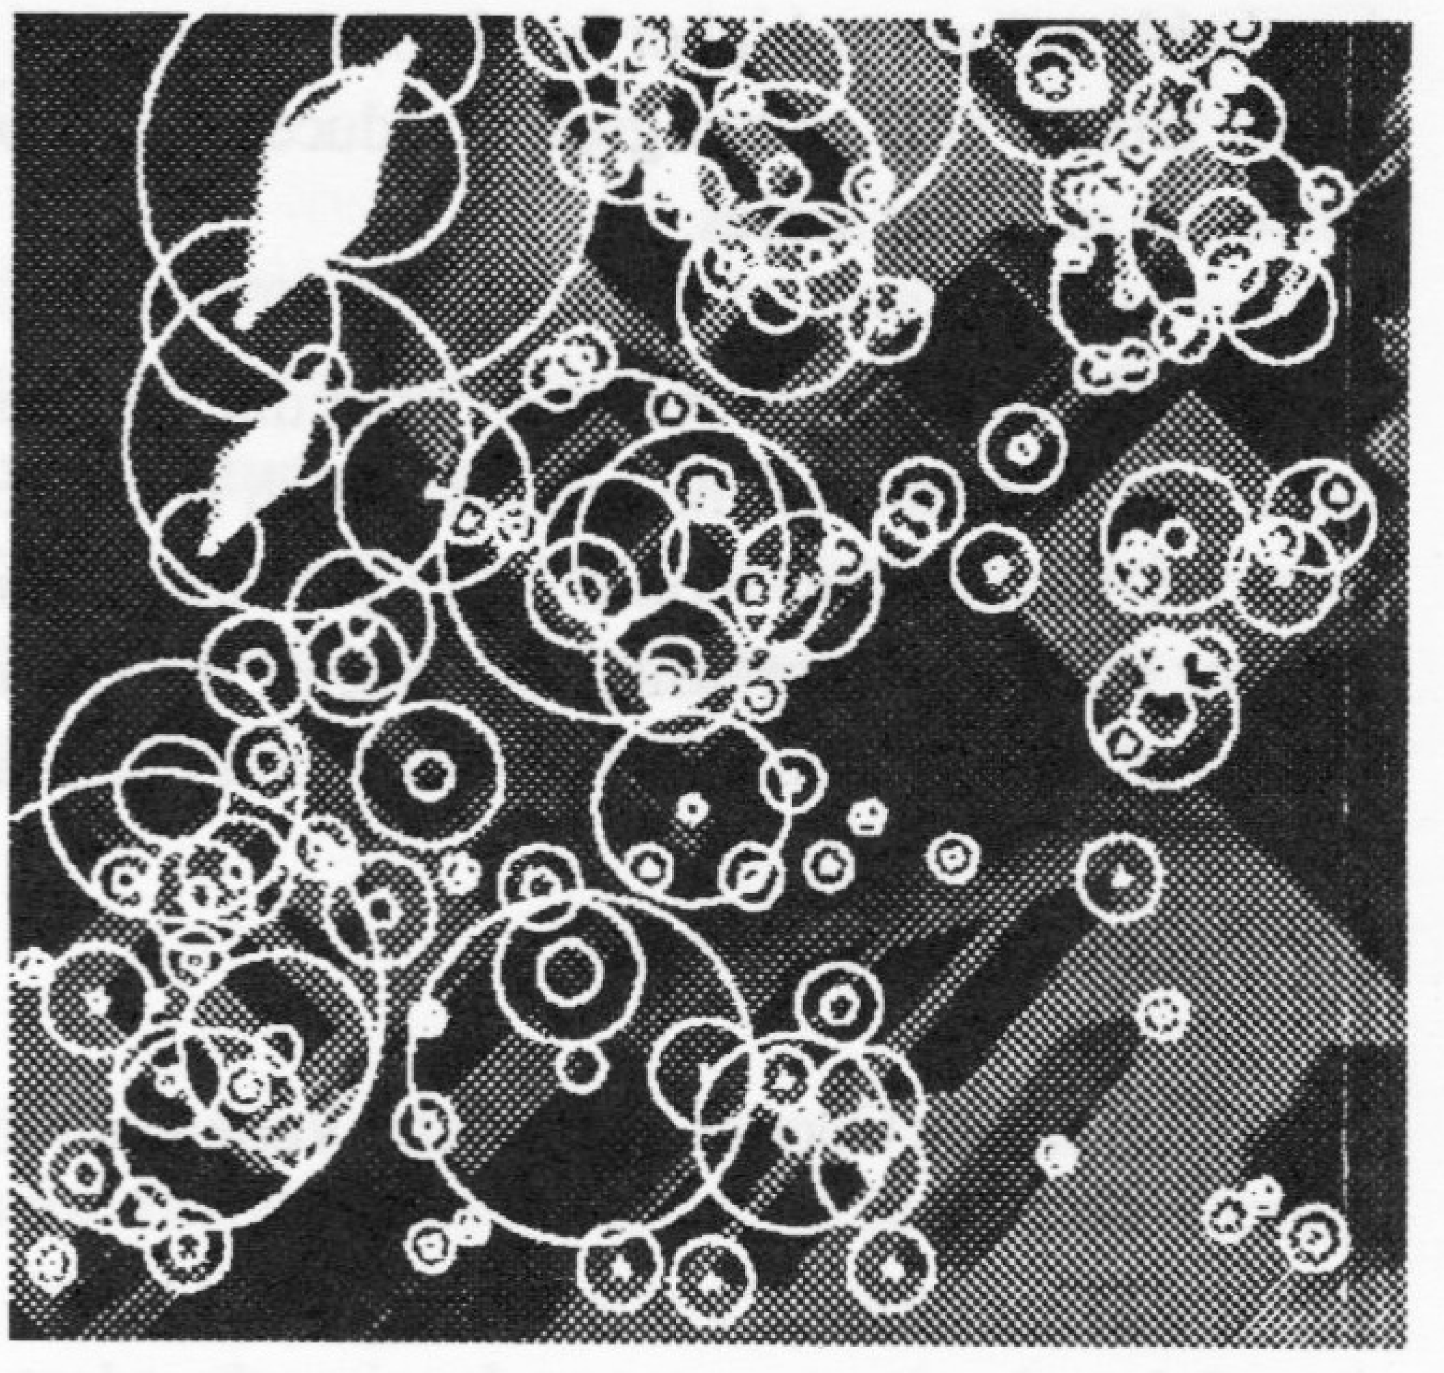

figure 7 - Multi-scale median detector

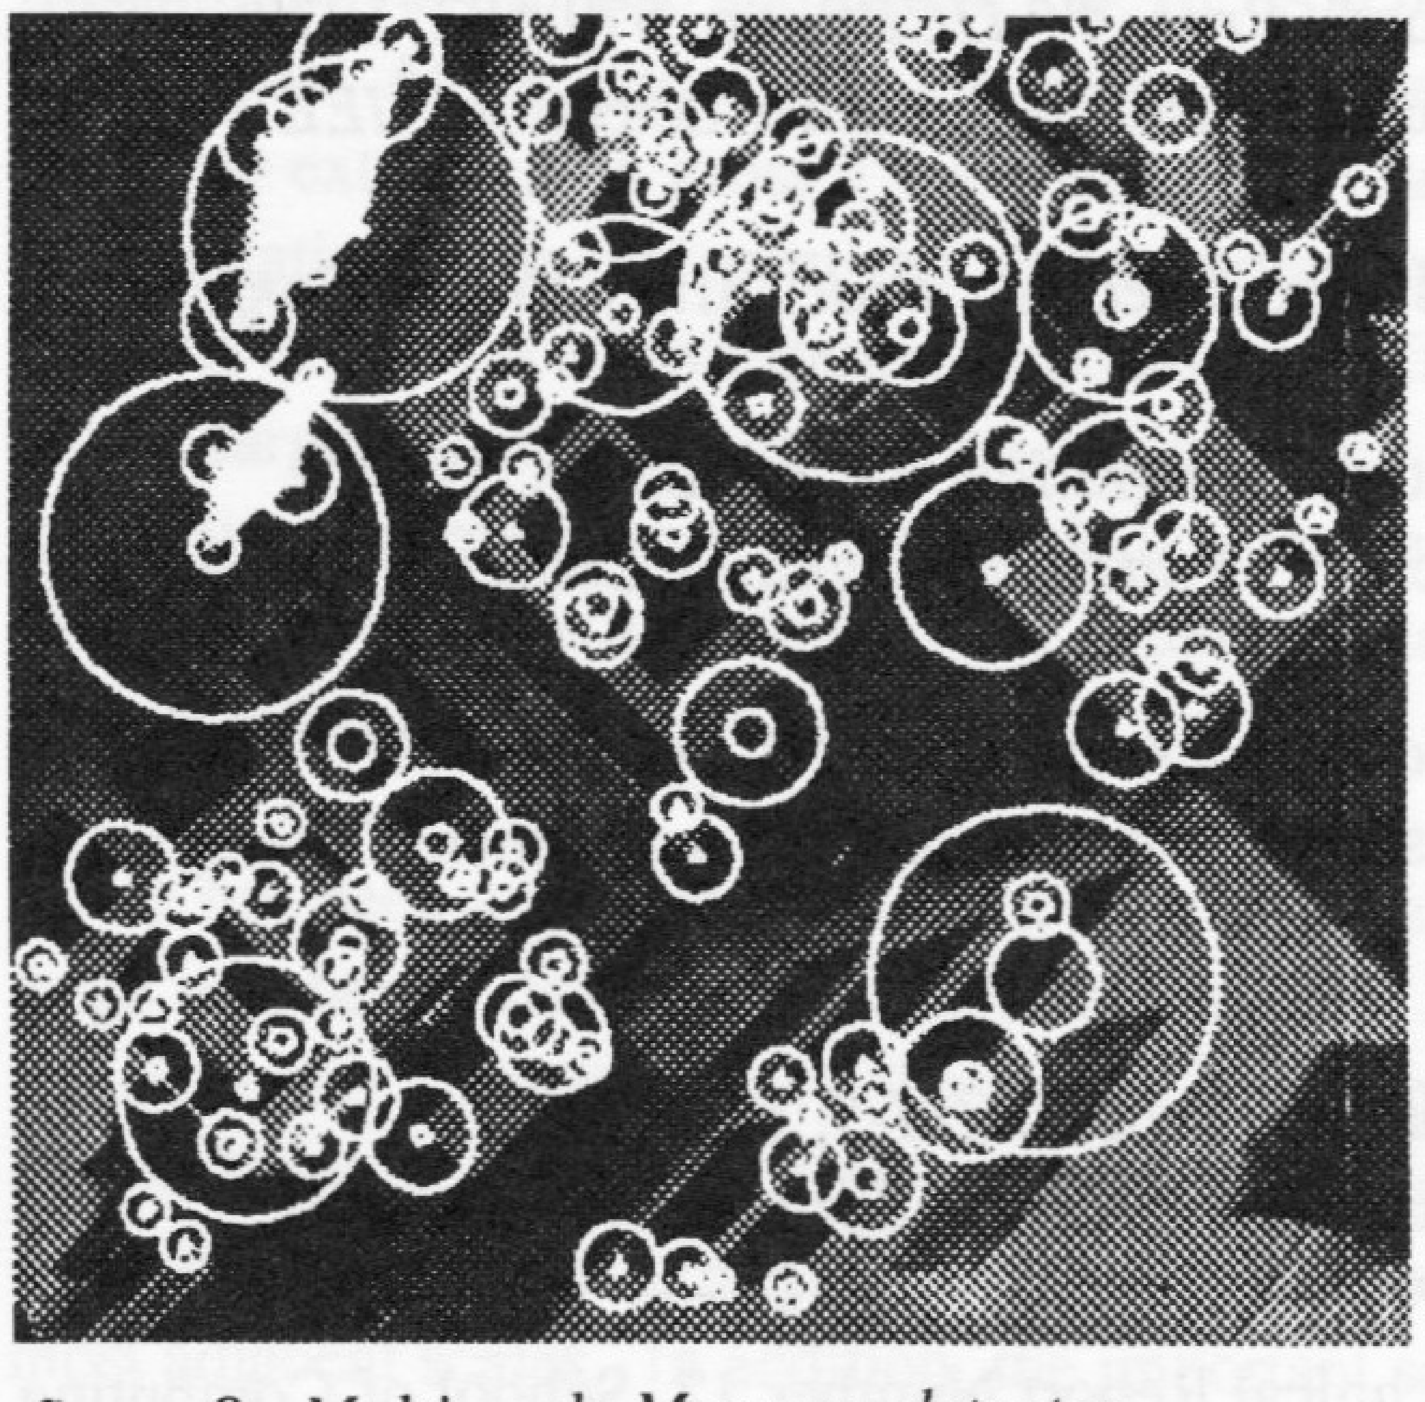

figure 8 - Multi-scale Moravec detector

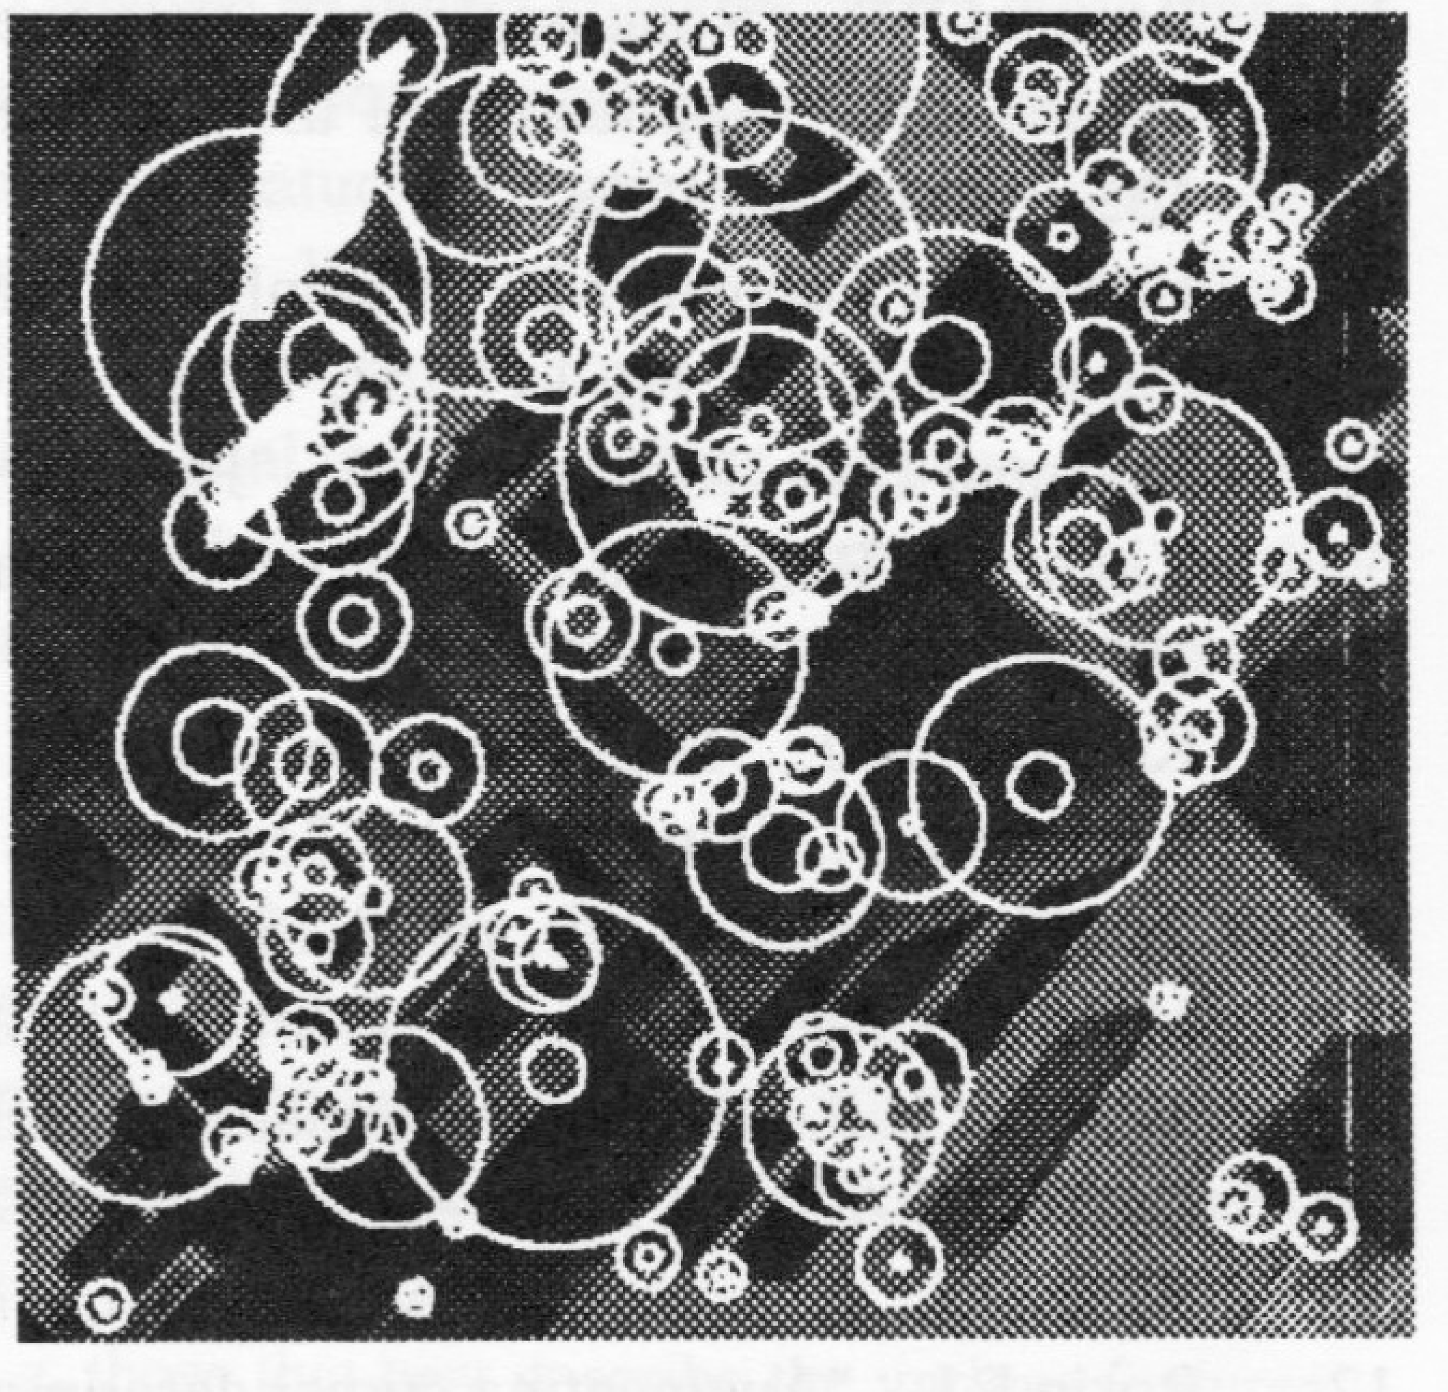

figure 9 - Multi-scale Plessey detector

5. Conclusions

This paper has asserted that standard grey-level corner detectors are inadequate since they restrict their attention to a single scale. Since corners in scenes are usually imperfect they cannot be reliable detected at any single scale, and a multi-scale analysis is required. It was shown that by applying a standard corner detector at multiple scales and then linking corresponding corners at different scales a variety of corner types likely to arise in scenes (i.e. noisy, textured, rounded, blurred) could be detected more successfully.

Several problems are evident in the multi-scale analysis of real images. In particular, in the process of blurring or subsampling the images to generate the grey-level pyramid topological changes may occur.

For instance, two close corners may be merged. This prevents corners of such objects being successfully tracked to the top end of their range of scales. Several solutions for preventing topological changes are possible. The image could be blurred by some adaptive technique that prevents blurring across regions (e.g. [13]). Alternatively, if the image could be successfully segmented, then the corners could be detected by analysing the curvature of the boundaries at multiple scales. A similar problem arises when nearby background clutter becomes included in the window that the corner detector operates within. This inevitably occurs at higher levels of the pyramid as the corner detector effectively operates over larger and larger areas. This clutter greatly reduces the response from the corner detector, again preventing successful detection at larger scales.

Finally, there is the problem of drift due to poor corner localisation by the corner detector and from the reduced resolution of the images which complicates the tracking of corners over scale. One approach to improve localisation would be to perform the multi-scale analysis over the complete image, analysing coarser scales by increasing the size of the detector's window rather than reducing the size of the image. However, this will be at the cost of greatly increased computational expense.

6. Acknowledgments

Thanks to James Cooper and Svetha Venkatesh for providing the code for several corner detectors.

7. References

1. Beaudet P.R., "Rotational invariant image operators", *Proc. 4th IJCPR*, pp. 579-583, 1978.
2. Davies E.R., "Application of the generalised Hough transform to corner detection", *IEE Proc. E*, **135**, pp. 49-54, 1988.
3. Harris C., Stephens M., "A combined corner and edge detector", *Proc. Alvey Vision Conference*, Manchester, UK, pp. 147-151, 1988.
4. Hoffman D.D., Richards W.A., "Parts of recognition", *Cognition*, **18**, pp. 65-96, 1984.
5. Kitchen L, Rosenfeld A., "Grey-level corner detection", *Pattern Recognition Letters*, **1**, pp. 95-102, 1982.
6. Koenderink J., "The structure of images", *Biological Cybernetics*, vol 50, pp. 363-370, 1984.
7. Lindeberg T., Eklundh J.-O., "Scale-space primal sketch: construction and experiments", *Image and Vision Computing*, **10**, pp. 3-18, 1992.
8. Marr D., *Computer Vision*, Freeman, 1982.
9. Moravec H.P., "Towards automatic obstacle avoidance", *Proc. IJCAI-5*, 1977.
10. Rosenfeld A. (ed.), *Multiresolution Image Processing and Analysis*, Springer-Verlag, 1984.
11. Rosin P.L., "Acquiring information from cues", *Proc. DICTA-91*, Melbourne, pp. 306-313, 1991.
12. Rosin P.L., "Augmenting corner descriptors", Technical Report Number 12, School of Computing Science, Curtin University, 1992.
13. Saint-Marc P., Chen J.-S., Medioni G., "Adaptive smoothing: A general tool for early vision", *IEEE Trans. PAMI*, **13**, pp. 514-529, 1991.
14. Sjöberg F., Bergholm F., "Extraction of diffuse edges by edge focussing", *Pattern Recognition Letters*, **7**, pp. 181-190, 1988.
15. Waltz D., "Understanding line drawings of scenes with shadows", in *The Psychology of Computer Vision*, ed. Winston P.H., McGraw-Hill, New York, 1975.
16. Whitten G., "Vertex space analysis and its application to model based object recognition", *Proc. CVPR*, pp. 847-857, 1988.
17. Witkin A.P., "Scale-Space Filtering", *Proc. 7th IJCAI*, pp. 1019-1022, 1983.