

An Analysis of Interest Operators for FIDO

Charles E. Thorpe

CMU-RI-TR-83-19

**Department of Computer Science and
The Robotics Institute
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213**

14 December 1983

Copyright © 1983 Carnegie-Mellon University

The CMU Rover has been supported at the Carnegie-Mellon University Robotics Institute since 1981 by the Office of Naval Research under contract number N00014-81-K-0503. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research or the U. S. Government.

Table of Contents

1. Introduction	2
2. The Moravec Interest Operator	4
3. Problems and Solutions	5
3.1 The Problem With Edges	5
3.1.1 Symptoms	5
3.1.2 Diagnosis	5
3.1.3 Prescription	5
3.2 The Problem With Scale	5
3.2.1 Symptoms	5
3.2.2 Diagnosis	5
3.2.3 Prescription	6
3.3 The Problem With 3-D	6
3.3.1 Symptoms	6
3.3.2 Diagnosis	7
3.3.3 Prescription	7
4. The New Interest Operator	8
5. Other Interest Operators	9
5.1 Kitchen-Rosenfeld / Nagel	9
5.2 Plane Fit	9
5.3 Auto-correlation	9
5.4 Controls	10
6. Experiments and Performance	11
7. Conclusions	13
7.1 Discussion	13
7.2 Summary	14
I. Test Results	15
Acknowledgements	17

Abstract

FIDO is a vision and navigation system for the CMU "Rover" mobile robot. The interest operator picks distinctive points to be tracked from image to image, both for stereo pairs of images and for images taken from a different Rover position. The performance of a simple interest operator is analyzed, and a theory formed to explain its defects. A new interest operator is built, implementing the suggested improvements. Tests run on the original and the new interest operators, as well as on several other operators, show almost identical performance. The reasons for the lack of improvement are discussed.

1. Introduction

An interest operator is a process that selects a few interesting points from one image to be located in one or more other images. The points selected should be easy to locate precisely in the other images. Points in bland areas are not very interesting, since all points look alike. Points on edges are somewhat more interesting, since their location is constrained in one dimension. But they are still not satisfactory because their exact location can slide anywhere along the edge. The points that are the easiest to locate, and thus the most interesting, are often corners of objects or high-contrast surface markings.

FIDO, the vision and navigation system for the Carnegie-Mellon "Rover" mobile robot, uses an interest operator as the first step of a stereo vision and visual motion determination process. When the Rover is standing still, it slides its camera from side to side and takes two or more pictures to form a stereo set. It runs an interest operator to pick about forty points from one of the images. FIDO then uses an image patch matching technique to find each point in each of the other images. Simple triangulation gives the three-dimensional position of each point relative to the Rover. The robot then moves forward, stops, and takes another stereo set of pictures. The points from the previous step are found in one of the new images by the same matching process, comparing patches around each interesting point in one of the old images with image patches in one of the new images. The points thus found are then stereo-ranged in the new set of images, and their new locations calculated. The apparent motion of the points from one Rover position to the next is used to calculate the actual robot motion.

There are several ways in which a better interest operator, one that picks more easily matchable points, will improve system performance. First, the more points that are tracked correctly, the easier it is to decide which points are tracked incorrectly and delete them. Secondly, more points and more precision can greatly improve the accuracy of the motion calculations [8]. Finally, the location of a point is only known relative to the Rover position at the time the point was picked. If the interest operator picks points that can be tracked for a longer time, their position will be known relative to an earlier, and presumably more accurate, robot position.

Early versions of FIDO used the Moravec interest operator, developed as a part of Moravec's thesis with the Stanford Cart [6]. The Moravec operator is simple and computationally efficient. But a close examination of its performance suggested that it could perhaps be improved. The points it picked were correctly matched about three quarters of the time; it seemed that a more sophisticated interest operator could perhaps be more successful in picking matchable points.

The research described in this paper set out to identify and correct the problems of the Moravec interest operator. Three specific shortcomings were found in Moravec's technique, and specific suggestions were made for improvement. This led to the design and implementation of a new (and hopefully improved) operator. Other interest operators, based on other approaches to the problem and from the literature, were implemented for comparison. As controls, points were picked both by a human expert and from a regular grid.

The various interest operators were evaluated by counting how many of the points they picked could be correctly matched by FIDO's matcher. The matcher uses a hierarchical correlation scheme that uses a pyramid of reduced resolution versions of each image. First, the general location of a point is found by

comparing a small patch around it in the lowest-resolution version of the source image with all possible locations in the lowest-resolution version of the target image. The location of the best match constrains the area that has to be searched in the next-to-lowest resolution image, and so on until the exact location of the point is found in the highest resolution version of the image. In each case, "best match" means highest pseudo-normalized correlation coefficient (see Appendix 6 in [6]).

This paper briefly describes the Moravec interest operator (Section 2), and discusses the problems with it and proposed solutions (Section 3). The new interest operator is presented in Section 4, and other operators in Section 5. Section 6 discusses the experiments comparing the various operators. Finally, Section 7 presents some analysis and conclusions from the work.

2. The Moravec Interest Operator

The Moravec interest operator measures the interest value of a point by taking the variance in four directions (horizontal, vertical, and the two main diagonals) of a small window (about 5 by 5 pixels) around that point and then choosing the minimum of those four variances. The directional variances are calculated by summing the squares of the differences in grey values of pixels adjacent along the direction currently being measured. For example, the following pseudocode fragment would calculate the horizontal directional variance.

```
for (row = minrow to maxrow)
  for (col = mincol to maxcol - 1)
    HorizontalVar += (image[row][col] - image[row][col+1]) ** 2;
```

Similar program fragments are used to calculate the vertical and two diagonal directional variances. The rationale for using variances is that bland areas will have low variance and therefore low interest value. The reason for using the minimum of the four directional variances is that points lying along a straight edge will have high variance across the edge, but low variance along the edge, and will still be uninteresting. Only points lying on corners or in textured areas will have a high response in all four directions. A local maximum filter checks each point against its 24 nearest neighbors to try to avoid bunching points in highly-textured areas.

The interest operator is applied to only one image in the pyramid of reduced-resolution versions, typically the next-to-highest resolution image. Part of the reason for not applying it to the original image is speed; every level up this pyramid corresponds to a speedup of 4. Another reason for not using the full-resolution image is that the reducing process, which averages pixel values, helps smooth out high-frequency noise.

3. Problems and Solutions

The Moravec interest operator does not always choose easily traceable points. Its failures can be attributed to three main causes: problems with edges at off angles, scale problems, and three-dimensional effects. This section presents the observed problems and their causes, and proposes some remedies.

3.1 The Problem With Edges

3.1.1 Symptoms

Even with the strategy of picking the minimum of the four directional variances, the interest operator still occasionally picks points on edges. This is especially disconcerting if a string of points along an edge is chosen, but the corners of the same object are not picked because of the local maximum filter. The points picked are rarely on horizontal, vertical, or 45-degree diagonal edges. Instead, they tend to lie on edges that lie somewhere between those directions.

3.1.2 Diagnosis

The major cause of this problem is the restriction of the variance calculations to the four orthogonal and diagonal directions. For instance, an edge oriented at 22.5 degrees will have high directional variance along the horizontal and 45 degree directions, as well as along the vertical and 135 degree directions.

3.1.3 Prescription

One solution would be to generate sums along all multiples of 22.5 degrees, instead of only multiples of 45 degrees. The next step would be to check at 11.25 degree intervals, and so on *ad infinitum*. A better approach is to find the direction of the gradient and to calculate the variance perpendicular to the gradient. This would, in effect, calculate the variance along an edge if an edge existed, no matter what the edge's orientation.

3.2 The Problem With Scale

3.2.1 Symptoms

Often the matching process, looking for an interesting point in a second image, will get lost in one of the low resolution images at the start of the search. If it misses the right point by enough of a margin in a low-resolution image, it can never recover. A related problem is that the matching process will sometimes get almost the right answer, then falter slightly at the last step.

3.2.2 Diagnosis

These problems are caused by running the interest operator at only one level of resolution. On the one hand, the distinctiveness of a point at a given level of the image hierarchy may be totally averaged out in the lower resolution images, giving nothing for the matcher to match. In the opposite case, something that looks very interesting may turn out to be, at the highest resolution, a bland patch. Then the matcher would find it correctly at all levels except the last.

An interesting footnote is that humans often make the same mistake. Experiments have been run with a human expert selecting points for the matching system to correlate. In cases where the matching failed, it was often because the point was easily detectable only at the highest resolutions. It is possible that people do the low-resolution matching so well that their only intuitive concern is for high precision at the last levels of matching.

3.2.3 Prescription

The points chosen by the interest operator must be interesting at several levels of resolution. The interest operator itself must therefore understand multiple levels of resolution. There are two obvious ways to do this: either use varying sizes of interest operators or use the same sized interest operator on multiple levels of the image pyramid. The latter is probably the right way to proceed. For one reason, program complexity is reduced by having a single version of the code. Also, running the same sized interest operator over successively smaller images means that the run time decreases significantly.

Another issue is how to combine the output from several levels of resolution into a single interest value at each point. Simply summing the interest value at each level of resolution is probably not adequate. It reduces, but does not eliminate, the problem of blandness at some levels of the hierarchy. An improved solution is to look at the minimum interest value over all levels of the hierarchy. Then, if a point has a high final value, it is guaranteed to be distinctive at all levels of the pyramid. One problem with this method is local maximum detection. A single pixel in a reduced image corresponds to many pixels in the highest-resolution image. If the minimum interest value for an area of the image is generated by one of the more reduced images, then several different locations in the original image will have the same interest value. Picking a local maximum is then impossible. A better solution would probably sort points according to their minimum value in the hierarchy, then sort according to sums of values when two or more points had the same minimum value.

A final issue is one of strictness. The stereo matcher uses, at each level, a search window twice as large as the area of the interesting patch. So the search can actually get lost by a few pixels at a given resolution and still recover the correct location at the next level. Requiring high interest value at the lower resolutions is then perhaps a bit of overkill. The present interest operator tries to ensure that a patch is distinct from its immediate neighbors. More sophisticated approaches would worry less about mistakes that could be corrected on the next level of the search, and concentrate instead on larger errors.

3.3 The Problem With 3-D

3.3.1 Symptoms

Sometimes the interest operator picks points that, while interesting at all levels of resolution and easy to find in the original picture, are still hard to locate in other pictures. Examples include contours of curved objects, "pseudocorners" formed by a distant object peering over the top of a nearer object, and highlights, which are the reflections of a light source. An extreme case is mentioned in Moravec's thesis, when the interest operator locked onto corners of shadows during an outdoor run. This would have been fine, except that the shadows changed as the sun moved, and the Cart got lost.

3.3.2 Diagnosis

All the cases above are examples of two-dimensional image phenomena that do not correspond to simple, viewer-independent, phenomena in the three-dimensional (3-D) scene. Good features, such as corners and surface markings, change appearance slowly and smoothly with small changes in lighting and in observer position. But many 3-D effects rely on coincidental alignments that change dramatically with small changes in the environment. The apparent location of the "edge" of a curved object, for example, shifts as the observer moves around it. Occlusions depend on alignment of the viewing position and the objects, and lighting phenomena depend on both viewer and light source locations.¹

3.3.3 Prescription

There is no easy solution to this problem. The interest operator works in the image domain, whereas the problems observed are in the scene domain. Incorporating 3-D scene knowledge into a vision program is difficult, slows down the program, and is not always successful. Furthermore, this is a chicken-and-egg problem. The interest operator is the first step in a stereo program that is designed to acquire 3-D data. If the detailed scene knowledge were available to add 3-D effects to the interest operator, there would be no need for the interest operator. The solution in FIDO's case is to track features from one robot position to the next and look for unstable features. Most of the features stay in a fixed relationship with each other from step to step. These are probably images of stable scene objects. Those features that wander are possibly caused by interactions between scene objects, lighting, and viewer position, and should be treated with suspicion.

¹This observation goes back at least as far as Leonardo Da Vinci, who said: "As regards all visible objects 3 things must be considered. These are the position of the eye which sees: that of the object seen [with regard] to the light, and the position of the light which illuminates the object." ([3], Section 115 "Of Painting")

4. The New Interest Operator

In response to the suggestions above, an improved interest operator would look at variance along the direction of edges and would be run at several levels of the image resolution hierarchy. A new interest operator was built to implement those recommendations. It has three main components.

1. A gradient operator, to find the direction (perpendicular to the maximum gradient) along which to take the variance. The new interest operator uses a Sobel gradient operator [2], either 3 or 5 pixels square.
2. A pixel interpolator. The direction along which the variance is taken need not be aligned with the rows, columns, or diagonals of the image. It is therefore necessary to get pixel values at non-integral locations. The pixel interpolation routine in the new interest operator uses a cubic approximation to the sinc function [9].
3. A variance calculator. This is a one-direction version of the variance calculator used by Moravec. It adds the squares of differences of the (interpolated) pixels adjacent along the edge, looking at all pixels in a small patch around each point.

The new interest operator calculates an interest value value at each point. It can be run at all levels of the image reduction pyramid. Other algorithms take the output interest values from the various levels, combine them, and do the local maximum detection.

The new interest operator typically calculates variance over a 5 by 5 window for each point. Each of the pixels in the 5 by 5 window is an interpolated value, and each interpolation requires the calculation of several cubics. Since a typical 256 by 240 image would require tens of millions of cubics, some sort of optimization would have to be done for this scheme to work in a real-time environment. One possibility would be to use some other operator to pick candidate points, then use the new operator to choose the best of the candidates.

5. Other Interest Operators

Several other interest operators were also implemented in the course of this project. This was partly to compare accuracy, and partly in the hope of finding a faster algorithm with nearly the same performance.

All interest operators were implemented to run at any level of the hierarchy. The result of a run was an interest value image the same size as the original, full-scale image. A post-processing program took any number of images and combined them into a single image, again the same size as the original. Two methods of combination were tried: averaging and finding the minimum at each pixel. Finally, another routine did the local maximum detection and picked the requested number of interesting points. This method of having each routine produce another image made it easy to see what kind of results were produced by each step of the processing.

5.1 Kitchen-Rosenfeld / Nagel

The first of the alternative interest operators is presented by Kitchen and Rosenfeld in a recent survey of corner-finders [4]. They found that this operator, which measures the turning of a fitted surface, was one of the best at isolating corners (section 2.4 of their paper; see also the conclusions in section 4). Nagel [7] has shown that this operator is equivalent to his own interest operator, which measures second order intensity variations. He also argues that this operator is more accurate than the corner-finder proposed by Marr and Hildreth [5], which finds the maxima of the curvature of zero-crossing contours of a symmetric second-derivative filter. (See also a paper by Canny [1] for another discussion of the limitations of the Marr-Hildreth operator at sharp corners.)

Kitchen and Rosenfeld give a computational complexity of $10n^2$ operations per evaluation of their operator, where n is the linear dimension of the operator. This can actually be reduced to a constant independent of the operator size by incrementally calculating the first and second derivatives of adjacent patches. Implementation speed need not be a concern therefore.

5.2 Plane Fit

A simple interest operator, called the "plane fit" operator, measures how poorly the pixel values match the best-fit plane. A plane can be fit to bland areas and ramp edges very well, to step edges less well, and only poorly fit to corners and textured areas. Then areas that are poorly fit are likely to be interesting points. So the plane fit operator fits a least-squares error plane to an image patch. It then sums the squares of residual error (distance from each pixel to the fitted plane). The higher the mean squared error, the more interesting the point.

5.3 Auto-correlation

A straightforward way to find points that are different from their neighbors is to look at auto-correlation functions. If a patch is bland, it will correlate very well with its neighbors. Patches that lie on an edge will correlate poorly as they are moved across the edge, but fairly well as they slide along the edge. Patches with interesting points will correlate poorly at all locations except for those that are exactly aligned.

Three auto-correlation operators were implemented. The first finds the highest correlation coefficient not at the origin. Intuitively, it calculates how much like this patch is the most similar nearby patch. The second version takes a simple average of correlation coefficients. The third takes a weighted average, with correlation coefficients closer to the origin weighted more than those further away.

Auto-correlations are especially appealing in the FIDO application. Since the point picked will be located in other images by a correlation process, this method appears to be the most direct way to select points that will have sharp peaks in their correlation functions.

5.4 Controls

Two other ways of picking interesting points were implemented. The first method picked points from a grid, spaced every 32 pixels, that evenly covered the image. The second used a human expert, picking the points he thought most likely to match well.

6. Experiments and Performance

The criterion for good performance of an interest operator is that the points picked be easy to locate in other images. Each interest operator was run under a variety of conditions to discover how well it met this criterion. The points selected were given to FIDO's point matcher, and a human observer counted the number of correct and incorrect matches in each case.

Two sets of pictures were used for testing. The first set is a pair of pictures of a telephone. The camera motion between the two pictures was not constrained to be in the focal plane. Epipolar lines are not known precisely. For the test runs, they were estimated to be within ± 12 pixels of horizontal, and the matcher was given that constraint. Since the camera motion included panning opposite the direction of translation, the image of a point may move either from right to left or vice versa, depending on its depth. These images have 8 bits of resolution, 256 rows by 256 columns.

The second set consists of three test images from one of the Stanford Cart runs. The scene is the inside of the old Stanford AI Lab; these images are referred to as SAIL images a, b, and c. The interest operators picked points in image a, which were then matched against both image b and image c. Images a and b are a stereo pair; the camera motion between them is known precisely, and is restricted to horizontal motion in the camera plane. Epipolar lines are still not known exactly, due to vidicon distortion that varies across the field of view. Again, for this test the matcher was told that motion was within ± 12 pixels of horizontal. Since there is no panning motion, there is a constraint on the horizontal location of a point: when the camera moves left, the image must move right. Two sets of runs were made: one with this horizontal location constraint and one without.

SAIL images a and c form a motion pair: they were both taken with the camera at its rightmost position, but at successive cart positions. There are no constraints on the camera motion between these two images. By matching image a against both images b and c, the tests measured both kinds of matching done during a real Rover run. These images have 6 bits of grey level, 240 rows by 256 columns.

Each interest operator was run on the right telephone image and on image a of the SAIL set. In each case, the operator was run over the original image, then over images that were reduced by factors of 2, 4, and 8. The four resulting images were combined two ways: by averaging and by taking the minimum value. Forty interesting points were extracted from each of four images: the full-resolution interest image, the interest image from the first reduced picture, and the two combined images. In all cases the interest operator size was 5 by 5, and the minimum separation between interesting points was 7 pixels. For the auto-correlation interest operators, the 5 by 5 patch was correlated over a 10 by 10 neighborhood.

The results were somewhat surprising. Among the Moravec, new, plane-fit, and Kitchen-Rosenfeld / Nagel operators, there was no significant difference in performance. Overall, all of them averaged between 30 and 33 correct matches out of 40 interesting points. Results were also similar with different ways of using the image hierarchy: the difference between the worst technique (picking interesting points from the first reduced image) and the best (from the minimum of all images) was on the average 1 extra point correctly matched out of 40. Detailed results are given in the Appendix.

The auto-correlation results ranged from slightly worse than the other methods to much worse. There were

some wide swings in performance with relatively small changes in parameters. The weighted average scheme, for instance, went from 6 to 23 correct matches when it was run on the first reduced telephone image rather than on the full-resolution image.

Points picked by the human expert were somewhat easier to match. In fact, the hand-picked points accounted for the only perfect matching of the whole data set. In all but one case, though, at least one of the interest operators was as good as or better than the human expert.

Points on a square grid were much harder to match. With all constraints, half of the points from the grid were correctly matched. With fewer constraints, the number correct fell below 30%.

7. Conclusions

7.1 Discussion

It is pertinent to ask why there was no improvement in matching results with the new interest operator. One reason is that the points selected by each method are in many cases the same. Of the 40 points chosen by the Moravec interest operator run on the full-resolution image, all but 11 were close to points selected by the new interest operator on the same image. All but 13 were also chosen by the Kitchen-Rosenfeld operator, and only 12 were not chosen by the plane fitting scheme. "Close" in this case means within 7 pixels, which is the minimum separation specified between points. There was also little difference among the different levels of resolution. For the Moravec operator, the points chosen from the full-resolution image were close to all but 7 of the half-resolution points, and close to all but 6 of the two sets of combination points. So to a large extent, the Moravec interest operator at a single resolution captures the same information as other more expensive schemes.

Another reason for the lack of improvement is matching errors due to 3-D effects. The new interest operator takes into account the first two problems mentioned in Chapter 3, but does not address the problems of 3-D. Although it is difficult to be sure about the cause of mismatches, spot checks seem to indicate that many errors are due to 3-D problems. These include:

- A large object in the foreground dominating small changes in the background. This is particularly apparent in the telephone pictures. Almost none of the interesting points on the wall are correctly matched, no matter which interest operator chooses them. While the lines between cinder blocks appear to be interesting, they all look like each other. Most matches get pulled to a rough position by the telephone in the foreground, which lines up the matcher incorrectly for the finer matching.
- Coincidental alignment. A good example of this is the large sloping object on the left hand of the SAIL images. Many interest operators picked a point where the right edge of that object intersects the door frame. This point appears to move up as the camera slides left, and to move down as the camera moves forward. It is always matched to the visually similar point in the next image, which is not in fact the same scene point.
- Reflections. The telephone has several highlights, especially on the right side of the handset. These shift from one image to the next as the relationship between the camera, object, and light source changes. Some shift small amounts (on highly curved surfaces), and are nearly the same scene point. Others shift further, and cause mismatches. This effect also shows up on the floor of the SAIL images. The floor has bright reflections of the lights at the far end of the ceiling. Those reflections shift from image to image, and can throw off the first levels of the matching process.

Some of the success of the human expert may be due to elimination of those 3-D problems. Humans are good enough at doing 3-D interpretation that they can certainly recognize and avoid pseudo-edges and coincidental alignments.

Another area of concern was the performance of the autocorrelation interest operators. They tended to pick lots of points in fairly bland areas, such as the floor in the SAIL pictures or the side of the telephone in the telephone pictures. One reason is that the correlation uses Moravec's pseudo-normalized correlation (see Appendix 6 of [6]). This normalizes for brightness of a patch by subtracting out the mean pixel values, and

partially normalizes for contrast differences by dividing by the sum of the squares of the pixel values. The net result is that small changes in a bland region will look nearly as interesting as larger changes in a higher-contrast image patch.

The poor performance of the auto-correlation interest operators can lead to speculation about the matcher. If normalization throws out so much information that the interest operator is ineffective, it might be that it also throws out information that would be useful to the matcher. Future work could certainly include unnormalized matchers and corresponding unnormalized interest operators.

The most dramatic improvement in the matcher's performance comes about not as a result of a better interest operator, but because of the horizontal hold constraint. Using this simple constraint decreased the average number of misses on the SAIL stereo pair (for new, Moravec, Kitchen, and plane-fit operators) from eight to three. This indicates that the most valuable use of real-time processing cycles is more likely to be in calculating and applying other such global constraints, rather than in trying to improve the interest operator.

7.2 Summary

- Regular grids of points do not work. There are significantly more mismatches with a regular grid than with any reasonable interest scheme.
- Cheap interest operators do work. Even the comparatively simple-minded operators, such as the plane fit or Moravec operators, are substantially better than regular grids.
- Expensive interest operators do not work any better than the cheap ones. The operators with carefully developed theories, such as the Kitchen-Rosenfeld and new interest operators, work about as well as the less expensive ones.
- Multiple levels of resolution do not contribute much. The points picked at a single level were about as easy to match as those that used several different resolutions.
- Autocorrelation interest operators do not work very well. Too much information is normalized out. Some of this could possibly be remedied by using a different (unnormalized) correlation coefficient.
- Other constraints are very important. Constraints on possible match location provide much more improvement than do the smartest interest operators. Even hand-picked points, with no constraints, were not as easy to match as automatically picked points matched with constraints.

I. Test Results

Interest op	Type	Epipolar	Horizontal hold	Motion	Telephone
New	0	31	37	34	24
	1	30	38	32	19
	min	31	36	34	20
	sum	33	38	34	18
Kitchen	0	29	38	37	21
	1	30	37	32	23
	min	33	38	34	27
	sum	31	37	33	21
Plane fit	0	34	35	30	25
	1	31	35	33	24
	min	32	36	33	19
	sum	32	36	31	24
Moravec	0	33	38	33	26
	1	30	37	32	27
	min	32	37	35	29
	sum	33	38	34	27
Auto Highest	0	7	11	7	18
	1	6	11	7	23
	min	23	26	27	16
	sum	17	22	21	14
Auto Average	0	31	35	31	6
	1	29	33	27	18
	min	27	31	27	12
	sum	30	33	27	12
Auto Weighted	0	6	8	4	6
	1	8	12	7	23
	min	22	27	26	12
	sum	26	29	29	11
Hand picked		34	40	35	26
Grid		14/49	24/49	21/49	16/49

All figures are number correctly matched out of 40 points, except for grid points which are number correct out of 49 points. Columns are:

- Interest operator name
- Which image were points take from. 0 for interest operator run on full resolution image, 1 for first reduced image, min for minimum over all 4 resolution images, and sum for average over all 4 images.

- Epipolar is the number correct (out of 40) on the SAIL stereo pair with no horizontal positioning constraint but with a maximum vertical change of 12 rows.
- Horizontal hold is the same as Epipolar, except the matcher used the constraint that points must appear to move right when the camera position moves left.
- Motion is the number correct (out of 40) for the SAIL motion pair, with no horizontal motion or epipolar constraints.
- Telephone is the number correct (out of 40) for the telephone pair. Like Epipolar, this set had no horizontal position constraint, but allowed for no more than 12 rows difference between the two images.

Acknowledgements

This paper owes a lot to Larry Matthies and Hans Moravec. Larry is the co-author of FIDO. He is also the "human expert" who hand-picked control points. Hans is the Principal Investigator of the CMU Rover project, as well as the author of the Moravec interest operator. He has supplied inspiration, high-level guidance, and advice at all hours of the night. I would also like to credit my advisor, Raj Reddy, for support and encouragement and for setting up and maintaining the excellent research environment of the Robotics Institute. Finally, thanks go to Cynthia Hibbard and Nancy Serviou for stylistic comments and suggestions.

References

- [1] J. Canny.
Finding Edges and Lines in Images.
Master's thesis, Massachusetts Institute of Technology, June, 1983.
- [2] Paul Cohen and Edward Feigenbaum.
The Handbook of Artificial Intelligence.
William Kaufman, 1982.
- [3] Leonardo Da Vinci.
The Notebooks of Leonardo Da Vinci.
Dover Publications, New York, 1970.
- [4] L. Kitchen and A. Rosenfeld.
Gray-level corner detection.
Pattern Recognition Letters 1:95-102, December, 1982.
- [5] D. Marr and E. Hildreth.
Theory of Edge Detection.
Proceedings of the Royal Society of London B 207:187-217, 1980.
- [6] Hans Moravec.
Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover.
Technical Report CMU-RI-TR-3, Carnegie-Mellon University Robotics Institute, September, 1980.
- [7] Hans-Helmut Nagel.
Displacement Vectors Derived from Second Order Intensity Variations in Image Sequences.
Mitteilung IfI-HH-M-97/82, IfI, March, 82.
- [8] Roger Tsai and Thomas Huang.
Analysis of 3-D Time Varying Scene.
Technical Report RC 9479 (#41904), IBM Watson Research Center, July, 82.
- [9] J. M. Vocar.
Image Magnification: Elementary with STARAN.
Technical Report GER-16342, Goodyear Aerospace Corporation, August, 1976.



Figure 1: Telephone Images



Figure 2: SAIL Stereo Pair



Figure 3: SAIL Motion Pair



Figure 4: Moravec Interest Operator. For all interest operator figures, top left picture is output when run on full-resolution image, top right is from half-resolution image, bottom left is quarter-resolution, and bottom right is output from operator run on image reduced by a factor of 8.

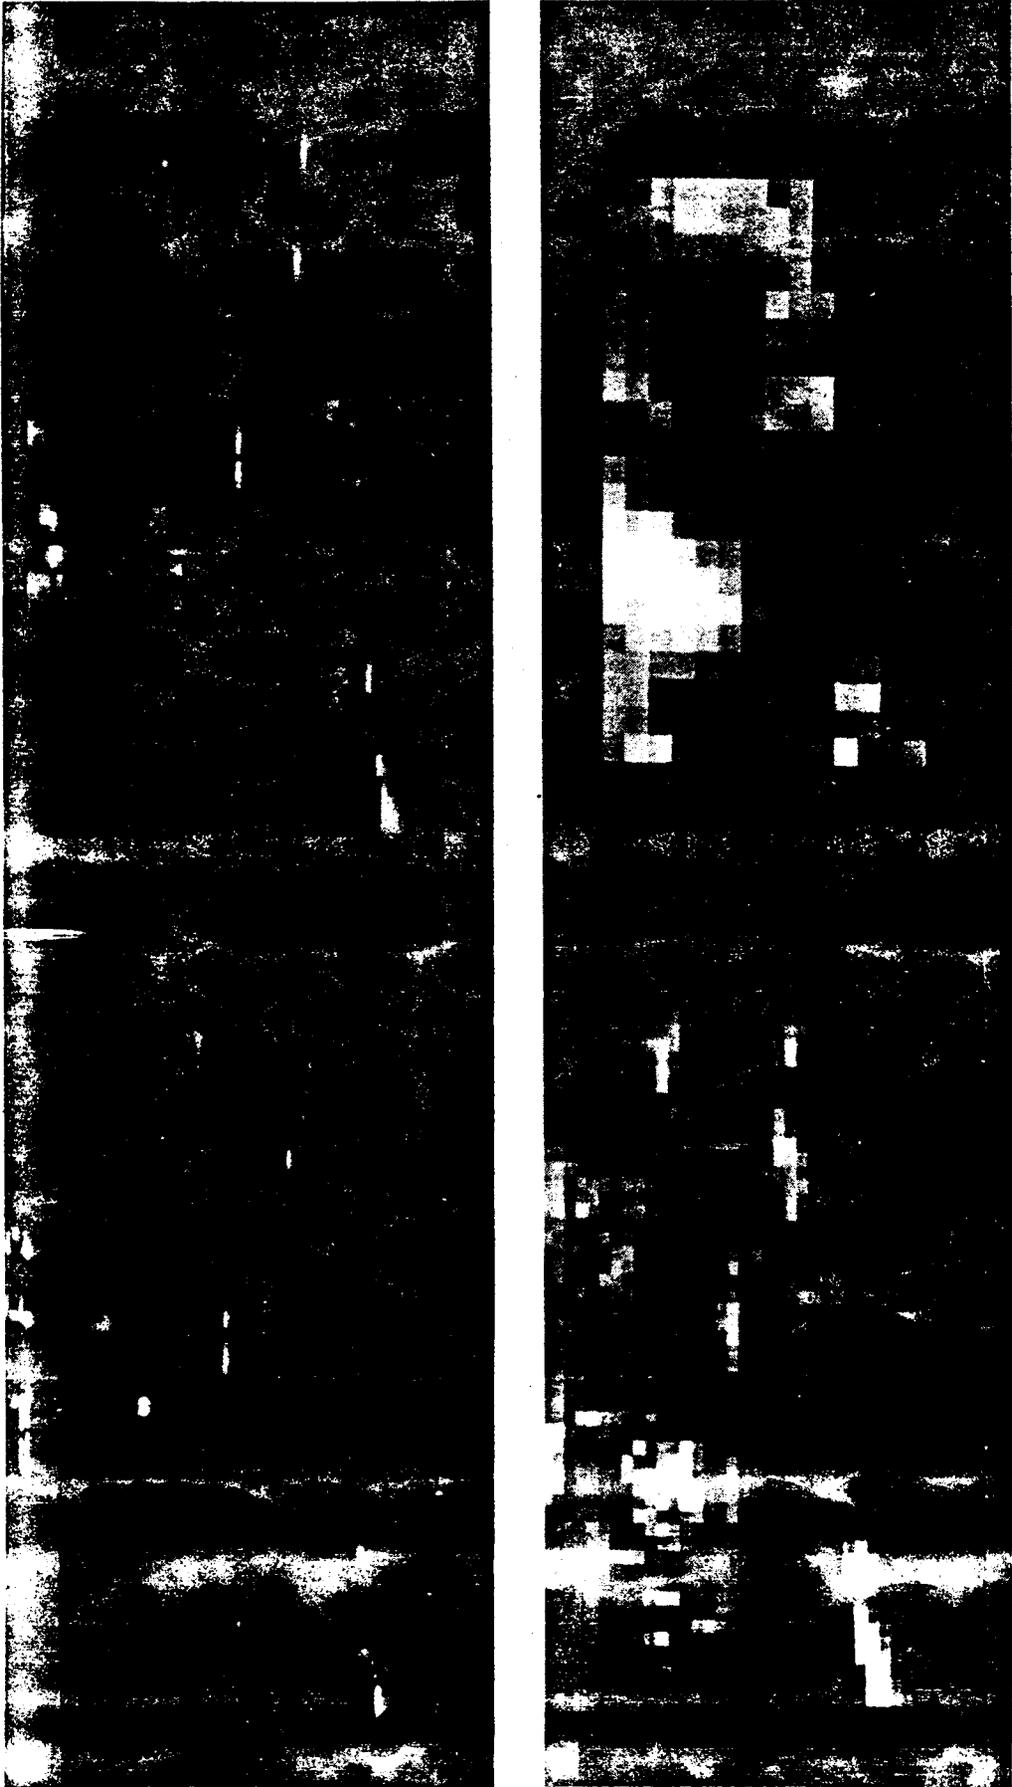


Figure 5: New Interest Operator

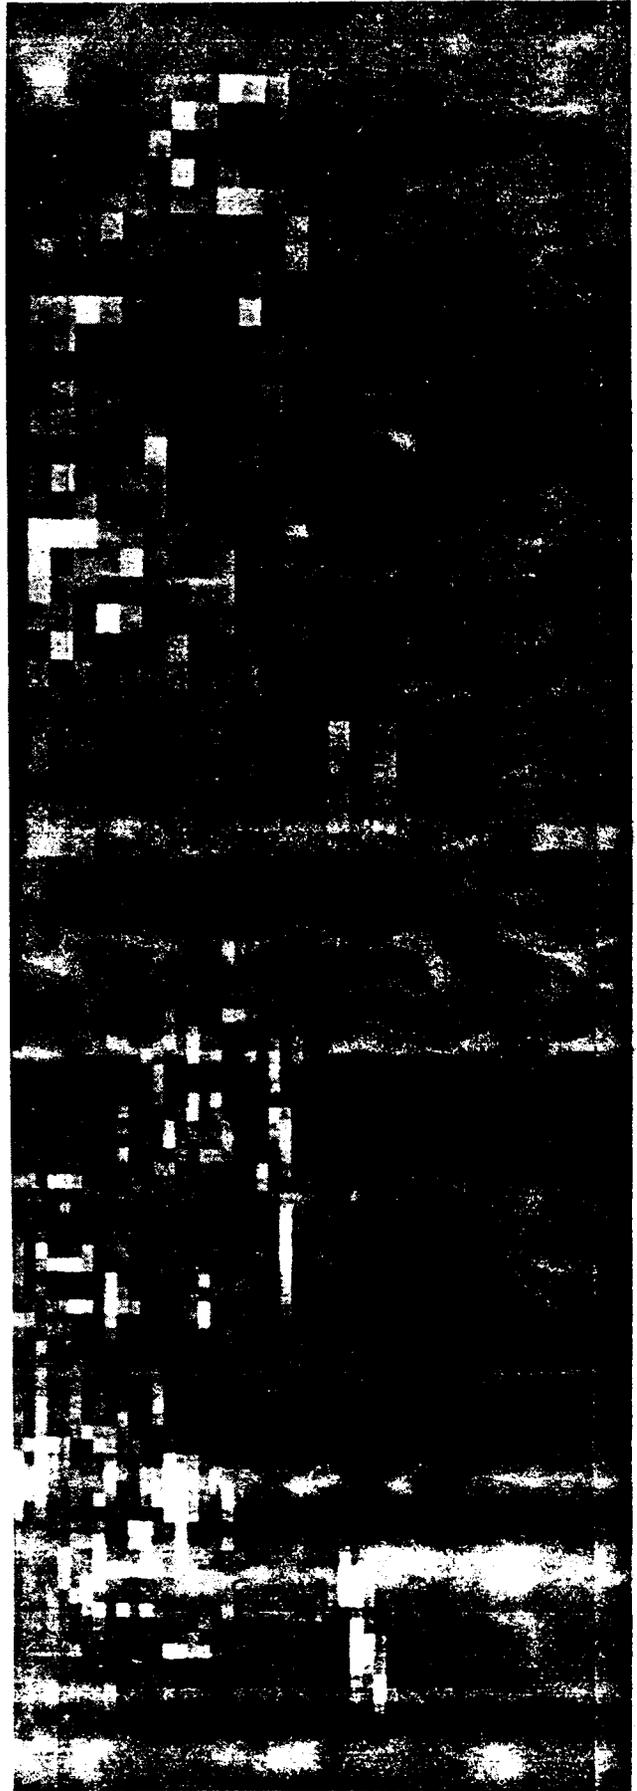


Figure 6: Kitchen-Rosenfeld / Nagel Interest Operator

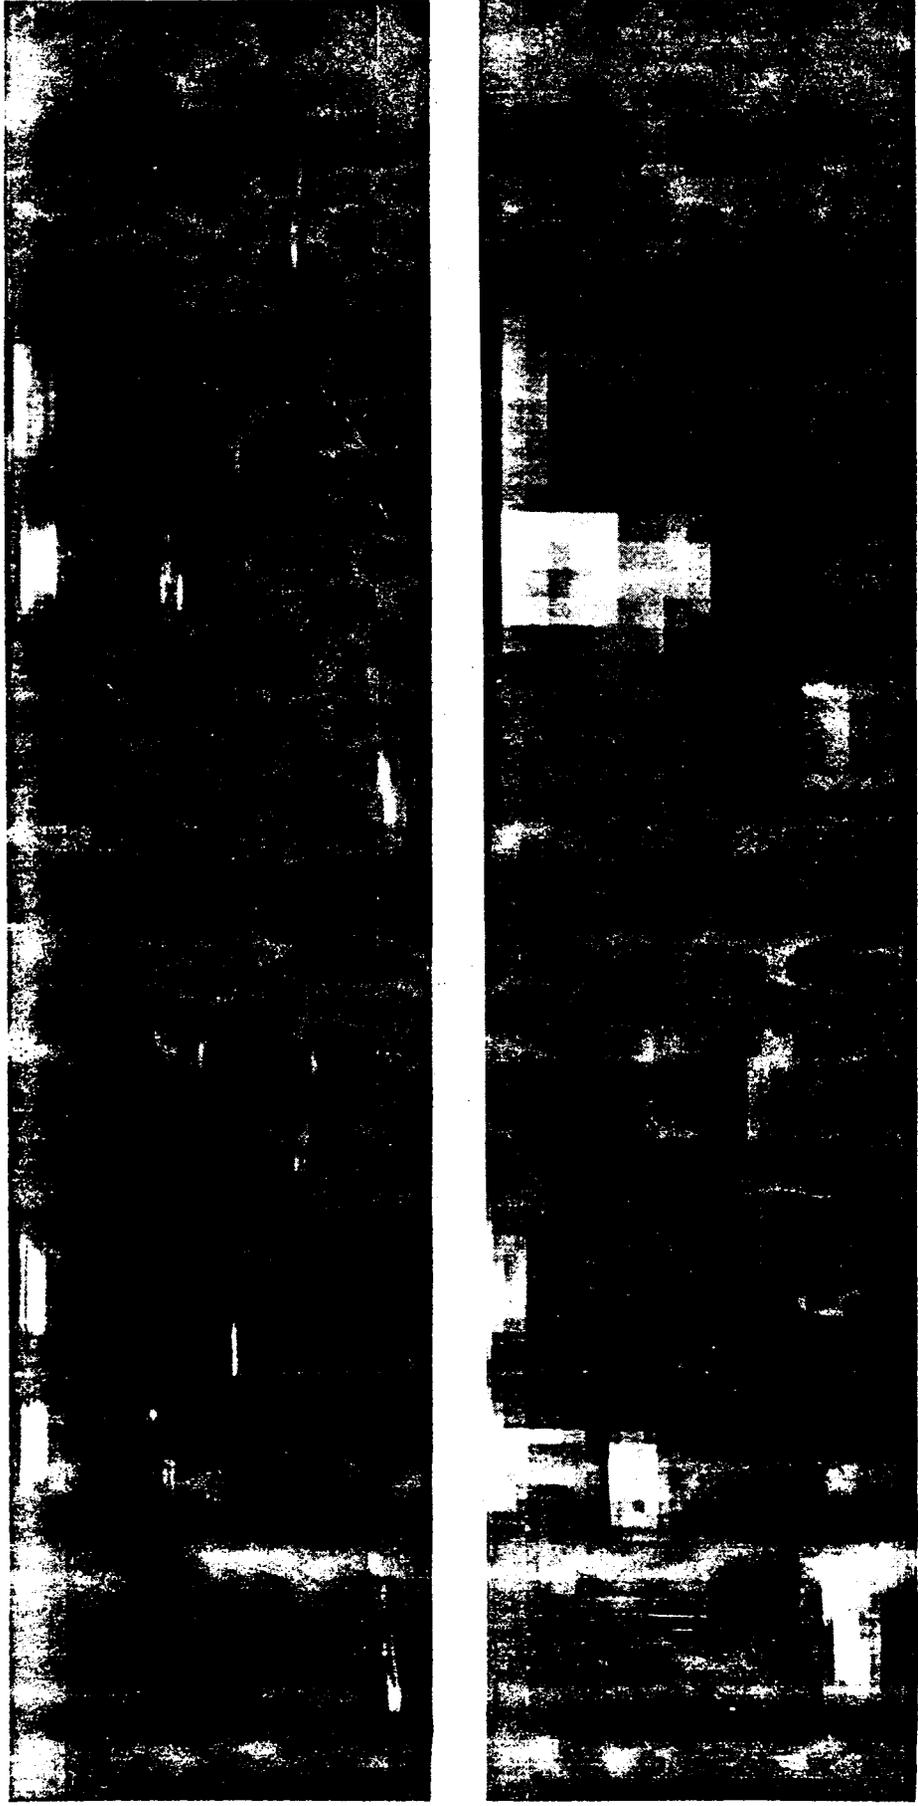


Figure 7: Plane-Fit Interest Operator

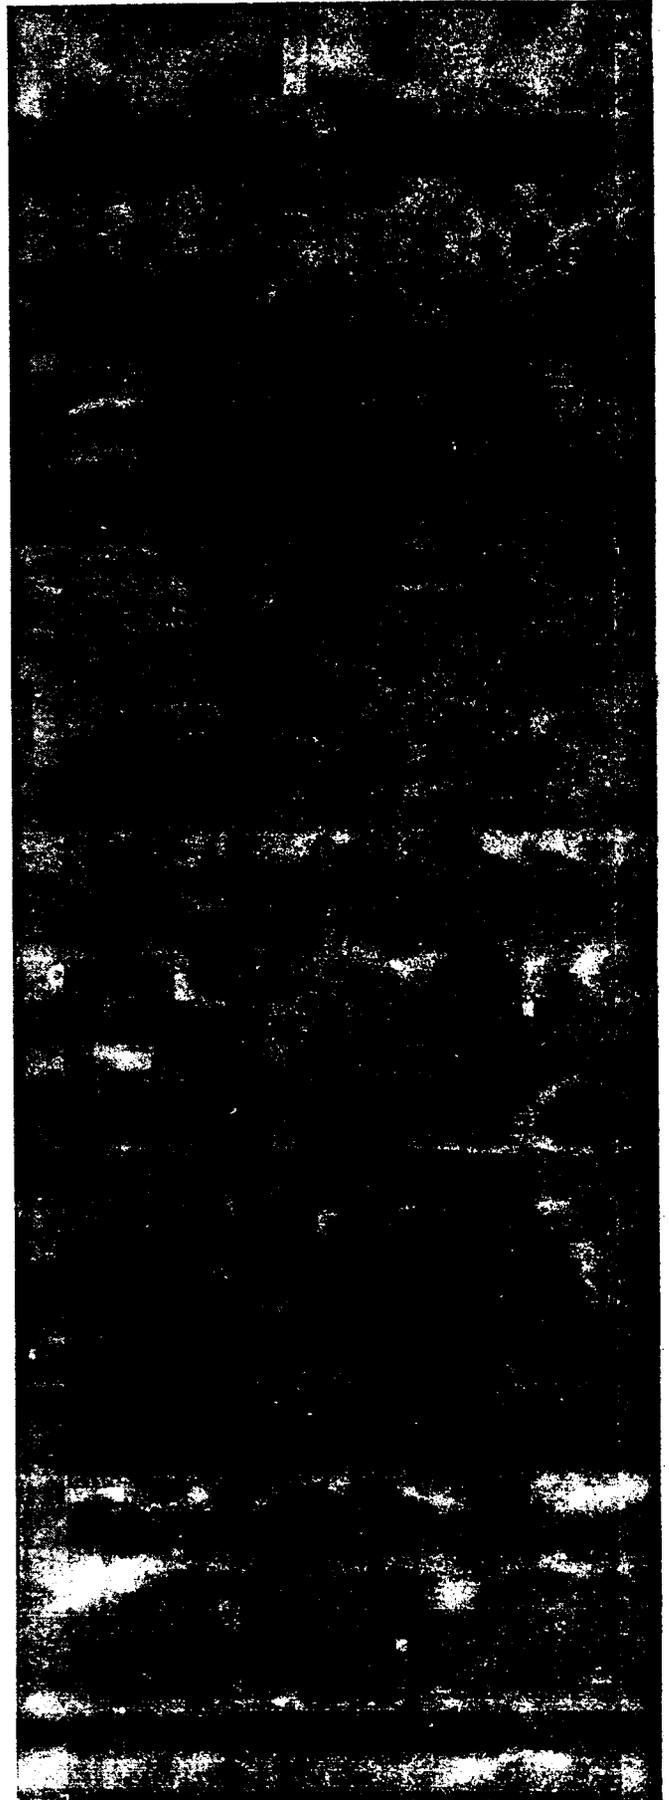


Figure 8: Auto-Highest Interest Operator

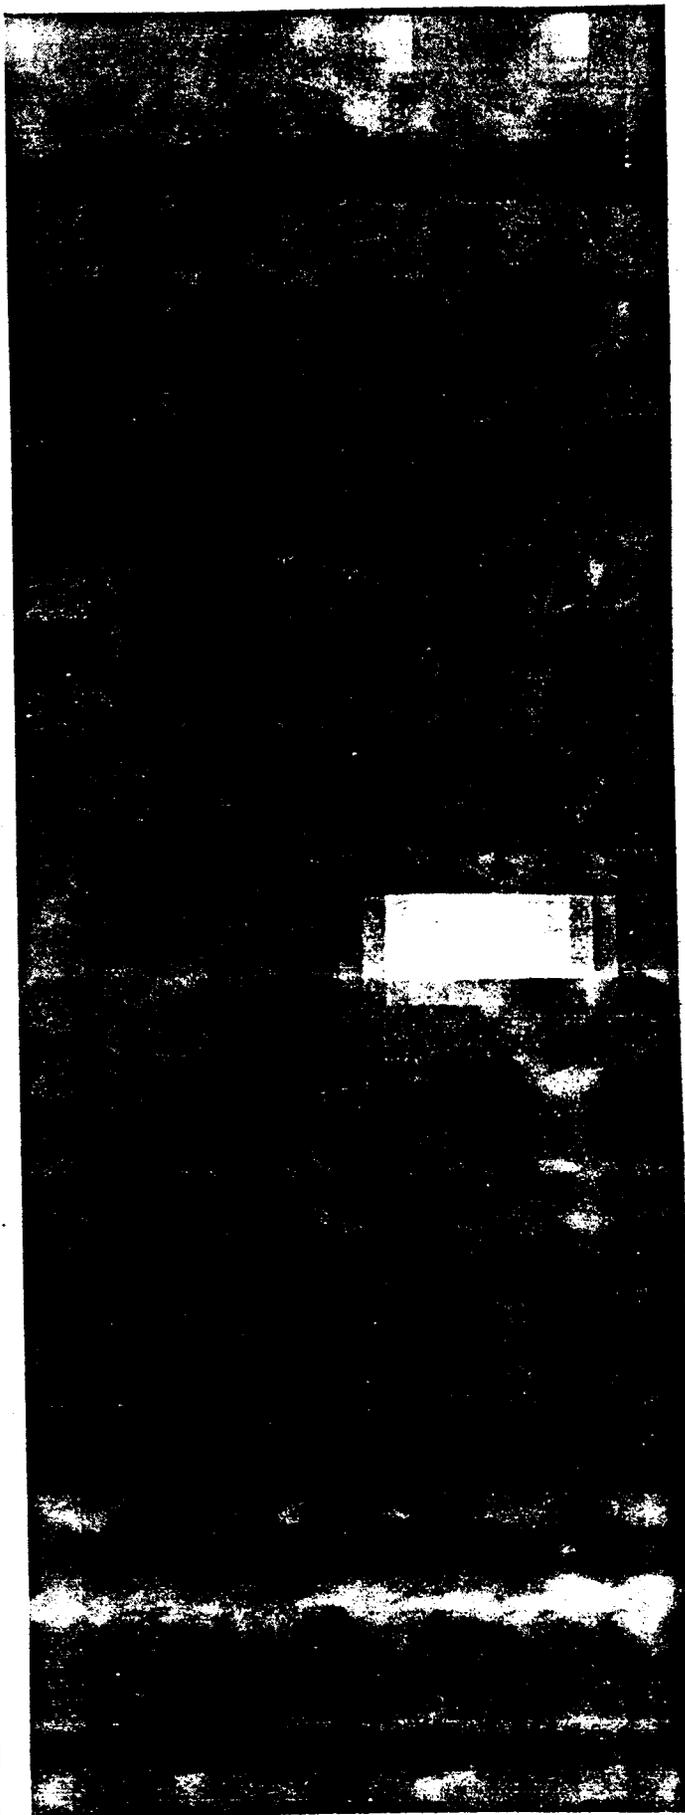
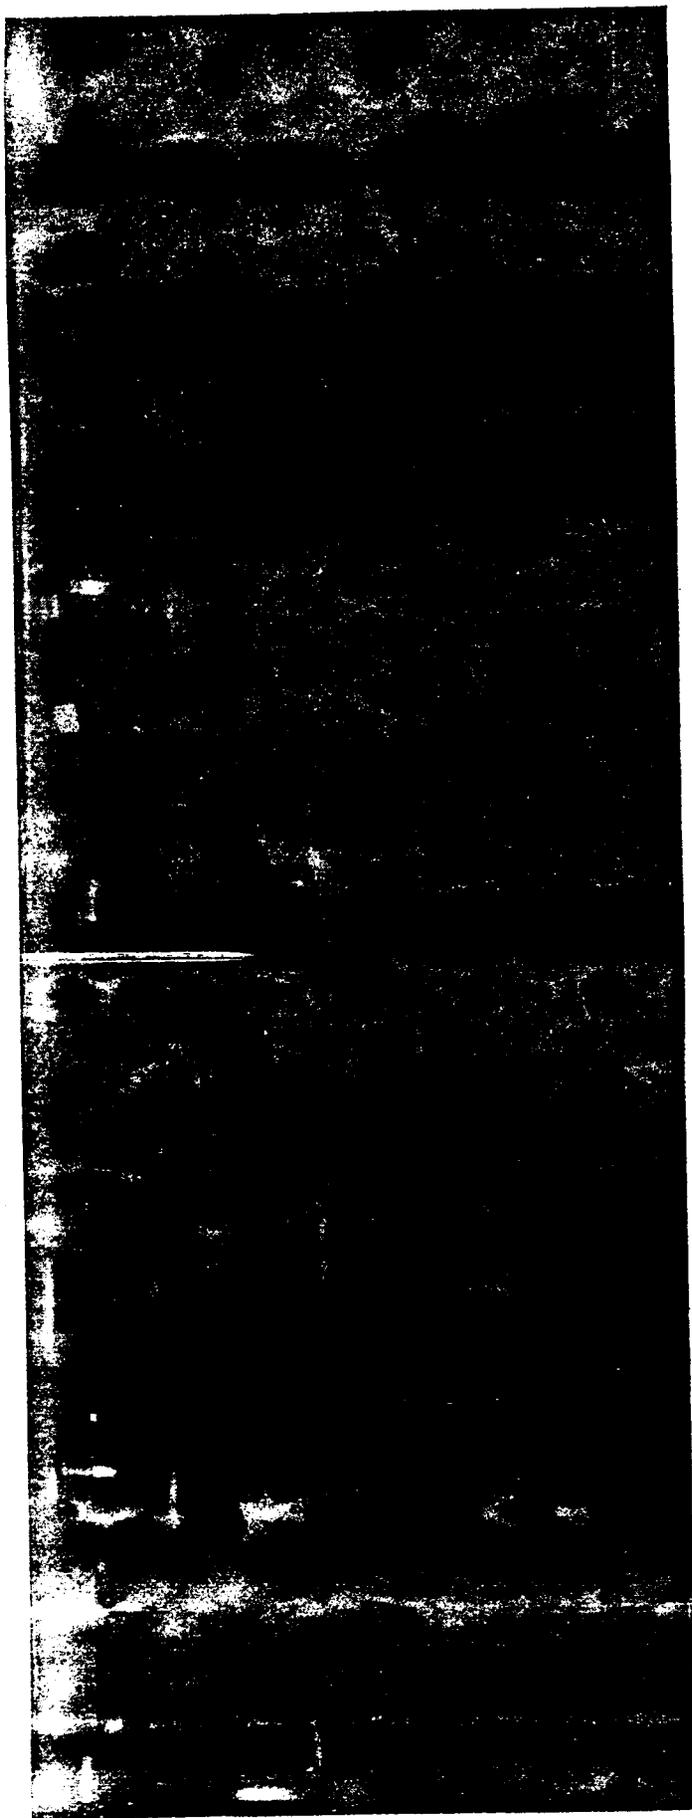


Figure 9: Auto-Average Interest Operator

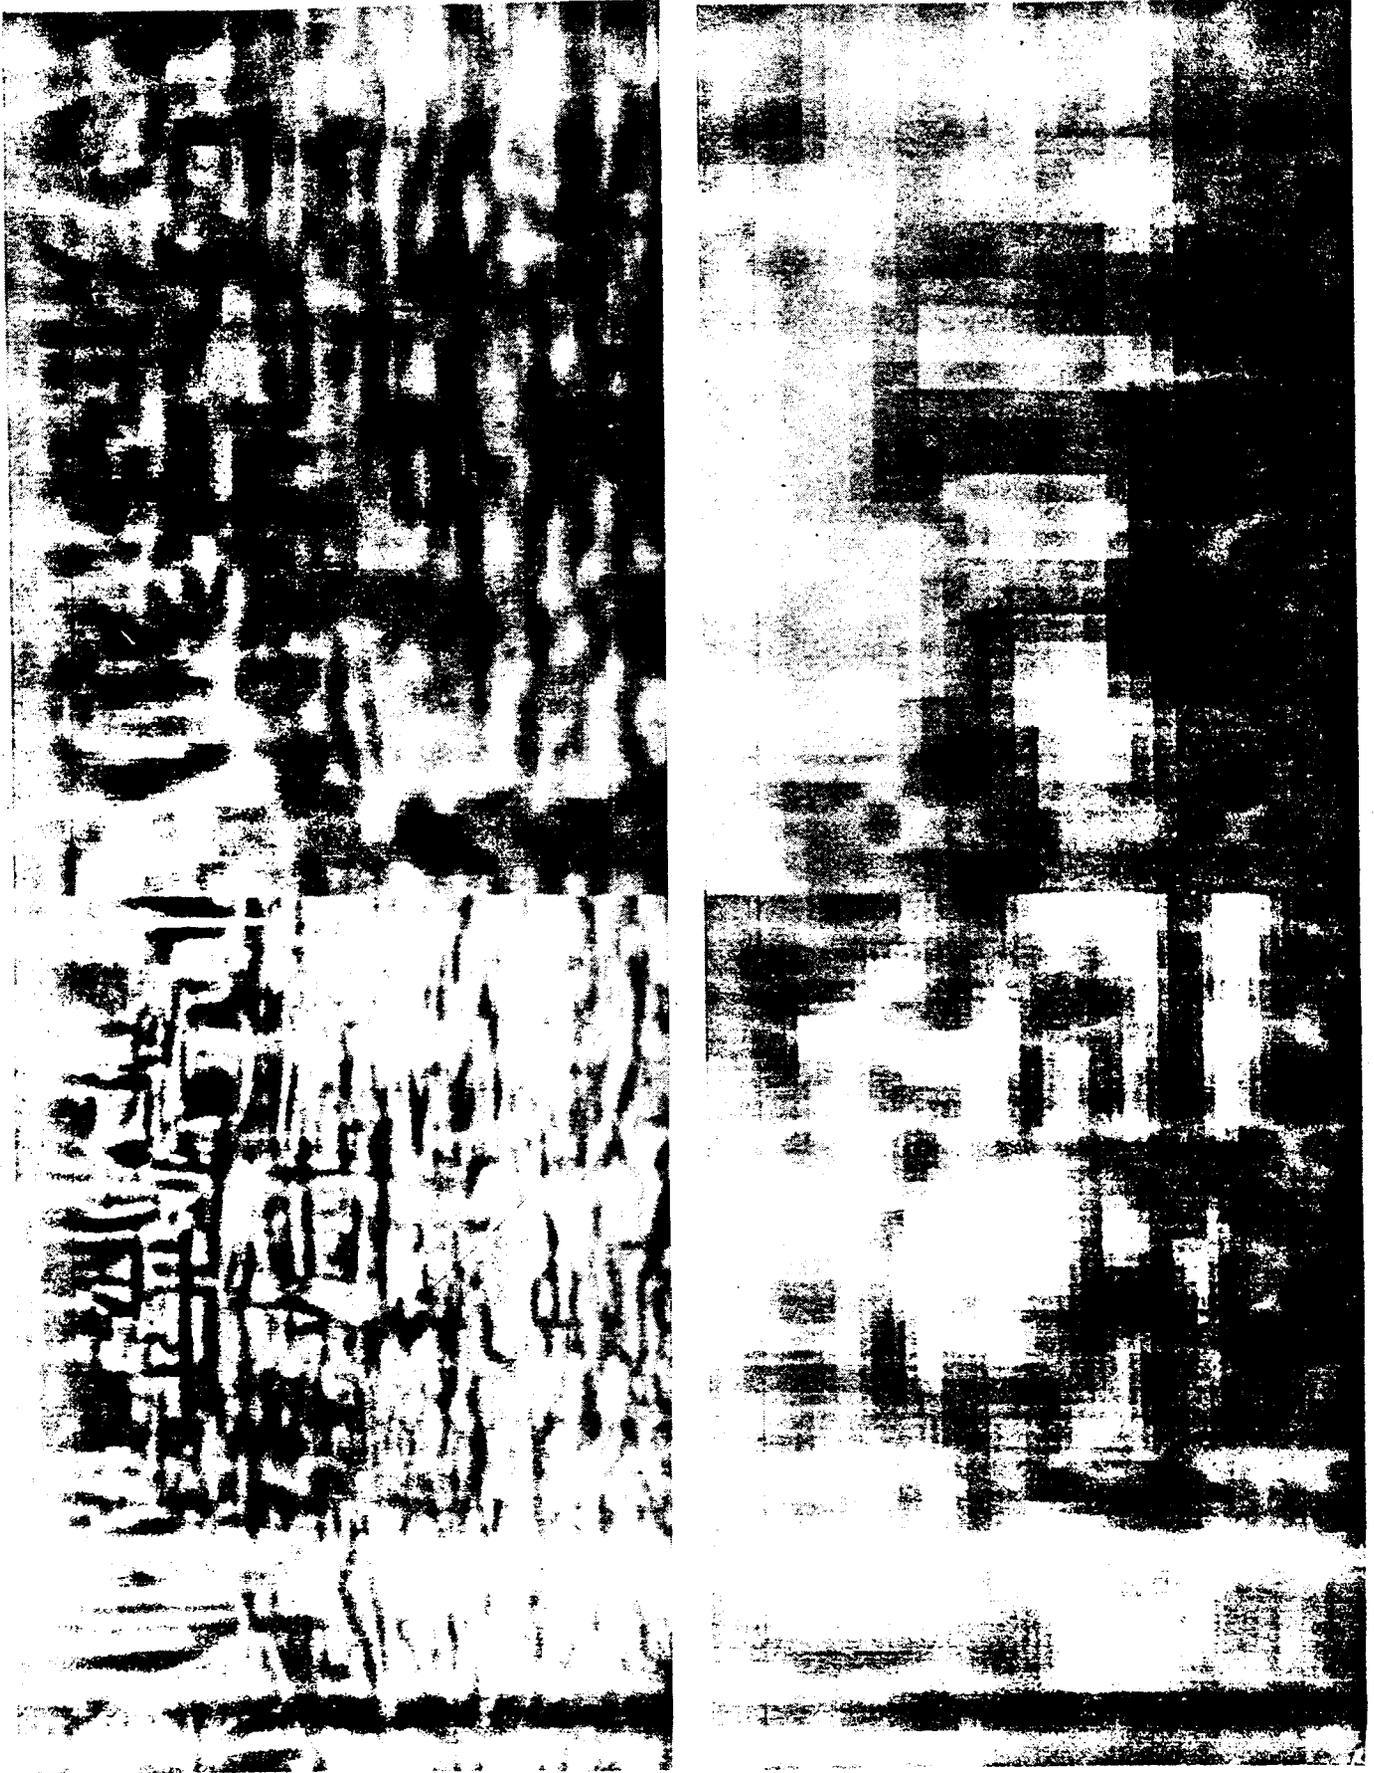


Figure 10: Auto-Weighted Interest Operator

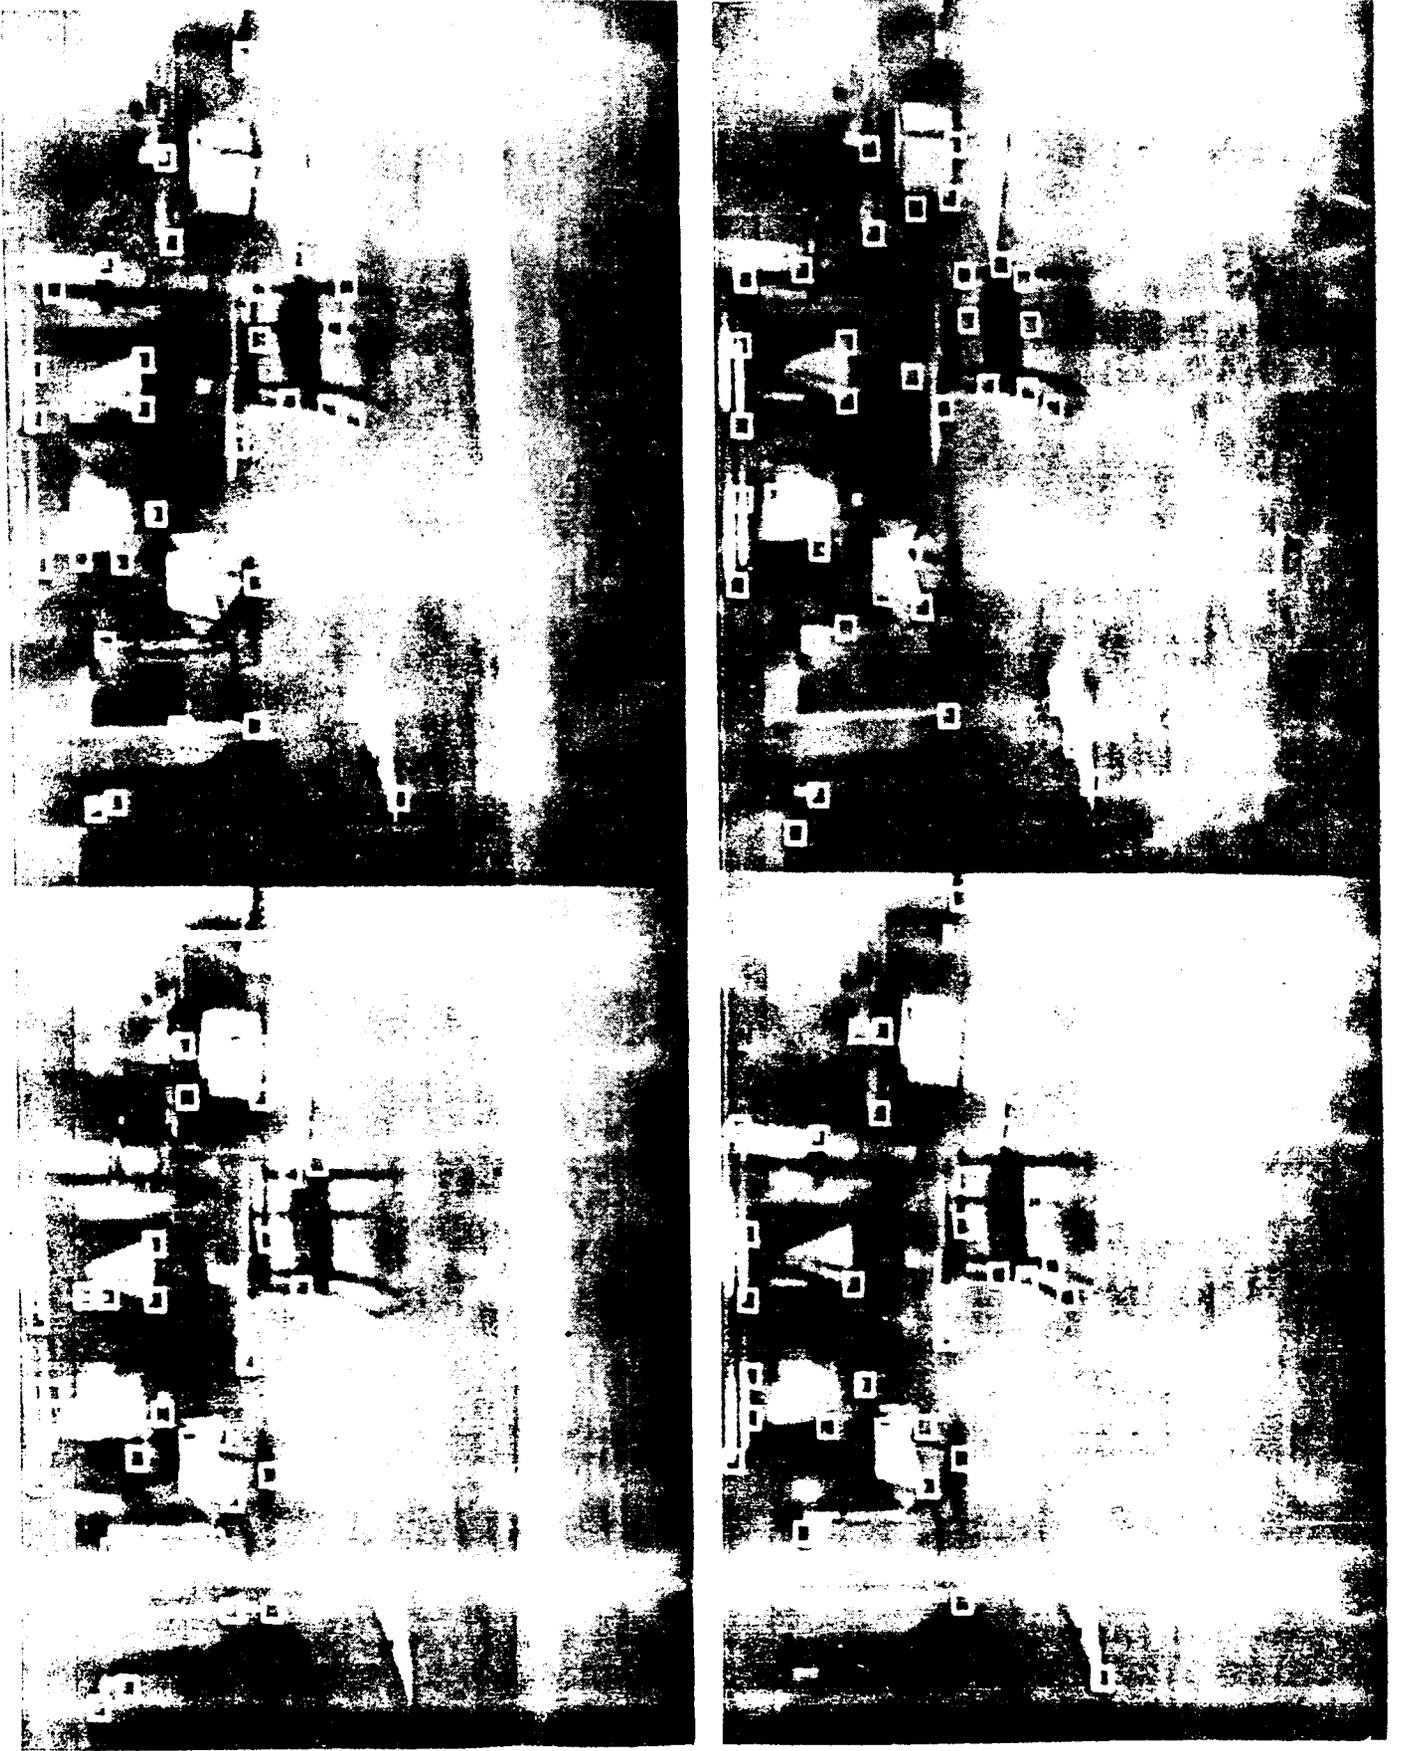


Figure 11: Points picked by New, Moravec, Plane-fit, Kitchen-Rosenfeld (clockwise from top left)

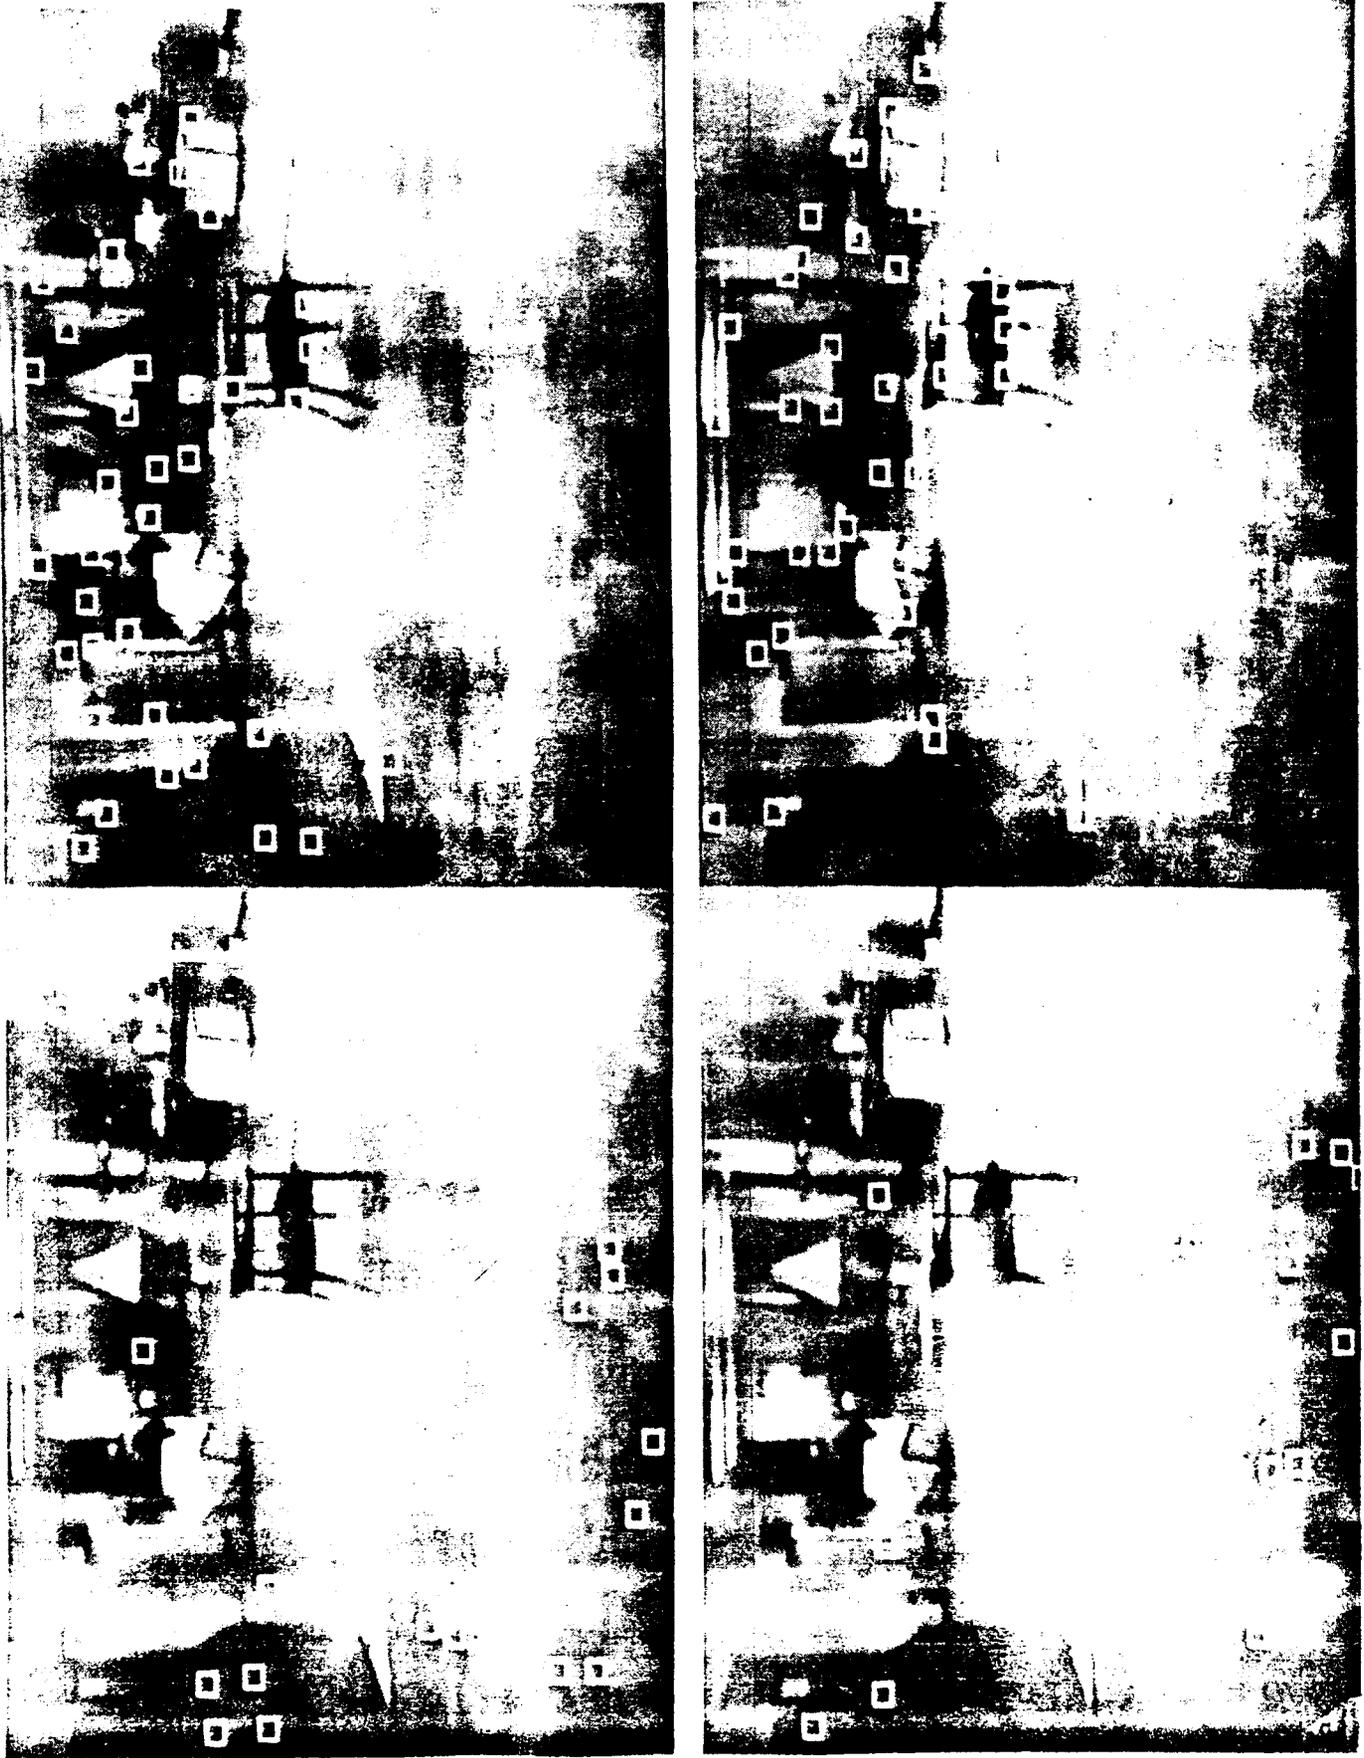


Figure 12: Points picked by Auto-Highest, Auto-Average, Human Expert, and Auto-Weighted (clockwise from top left)

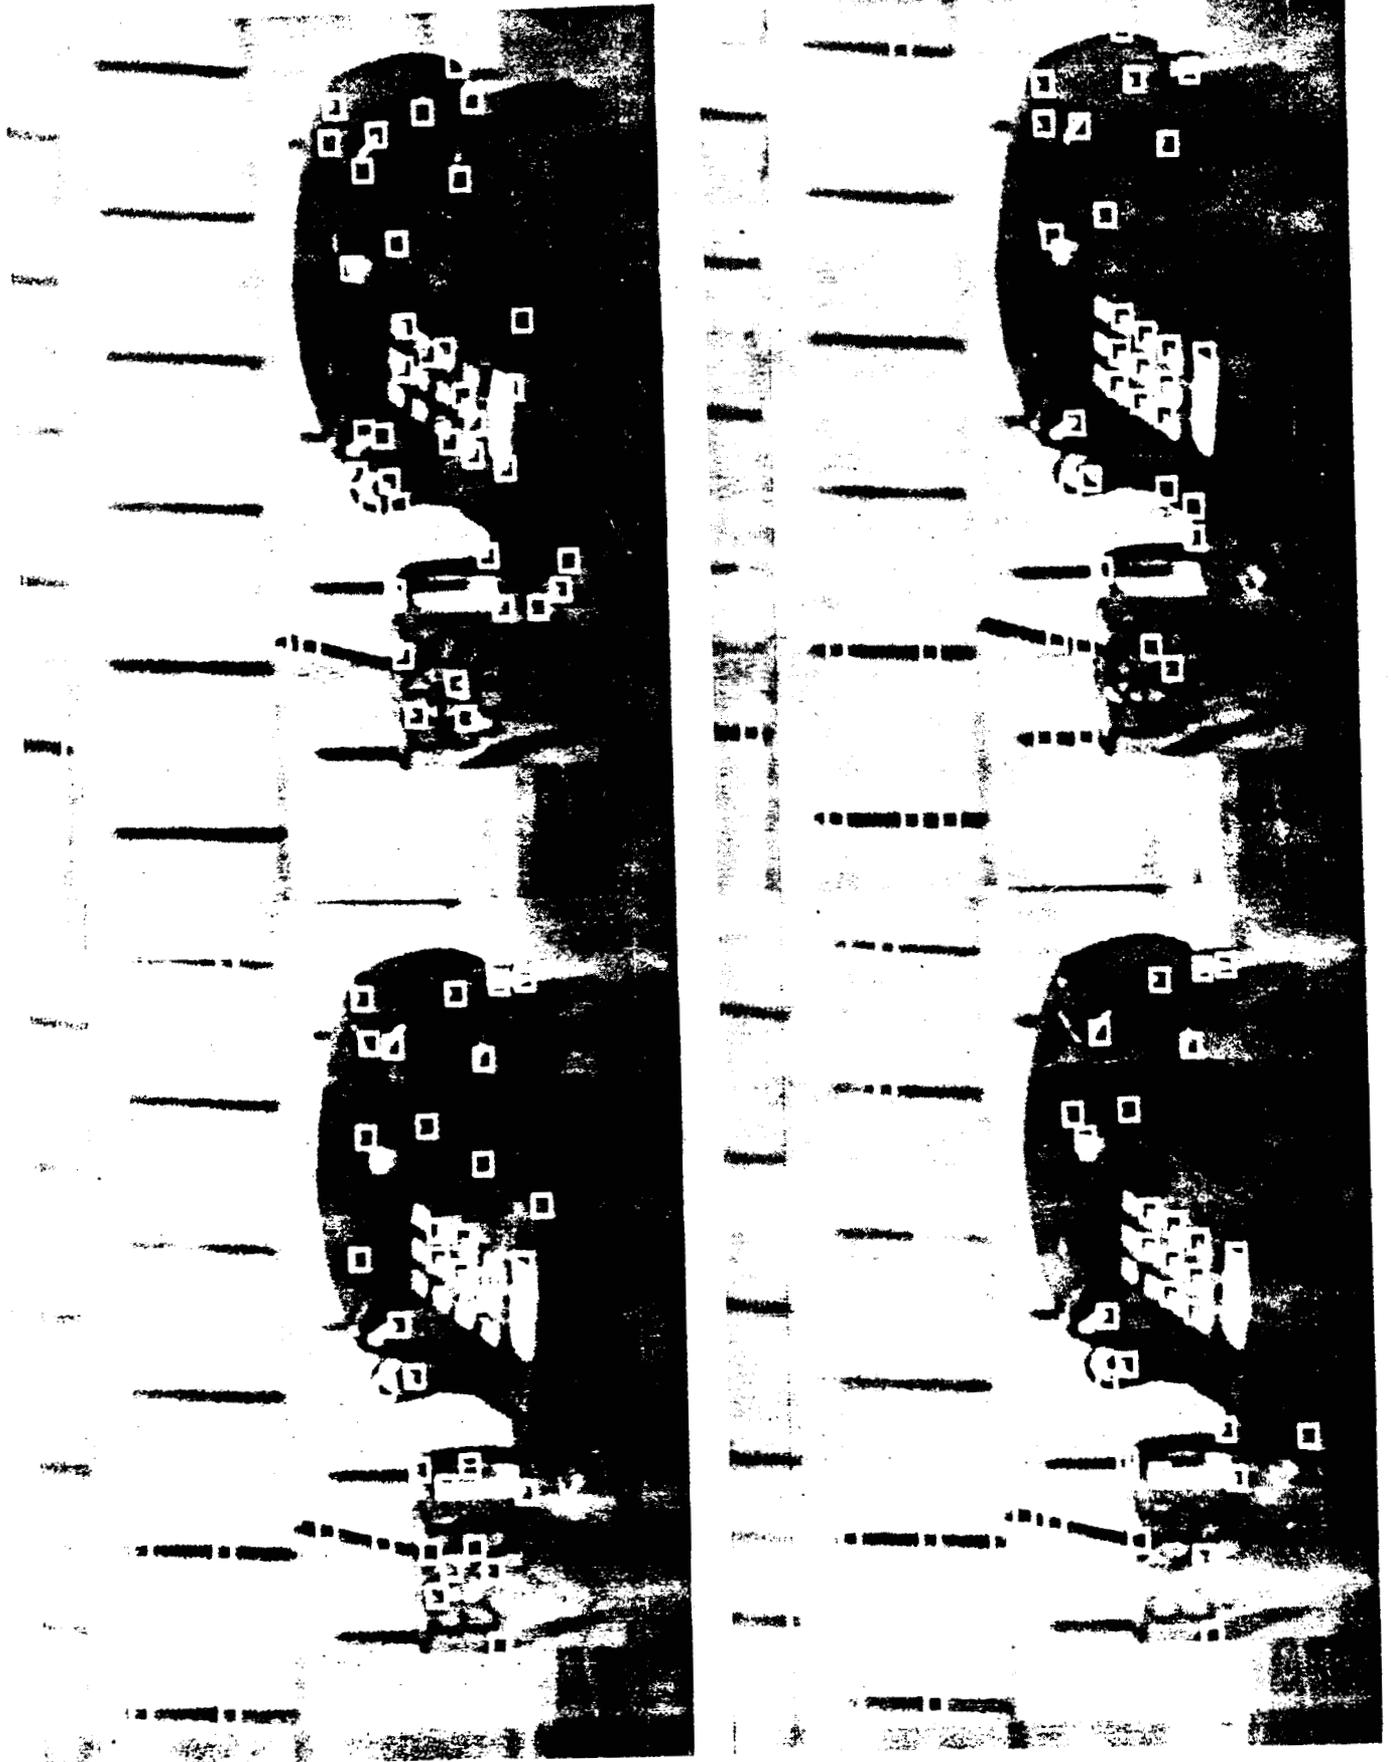


Figure 13: Points picked by New, Moravec, Plane-fit, Kitchen-Rosenfeld (clockwise from top left)

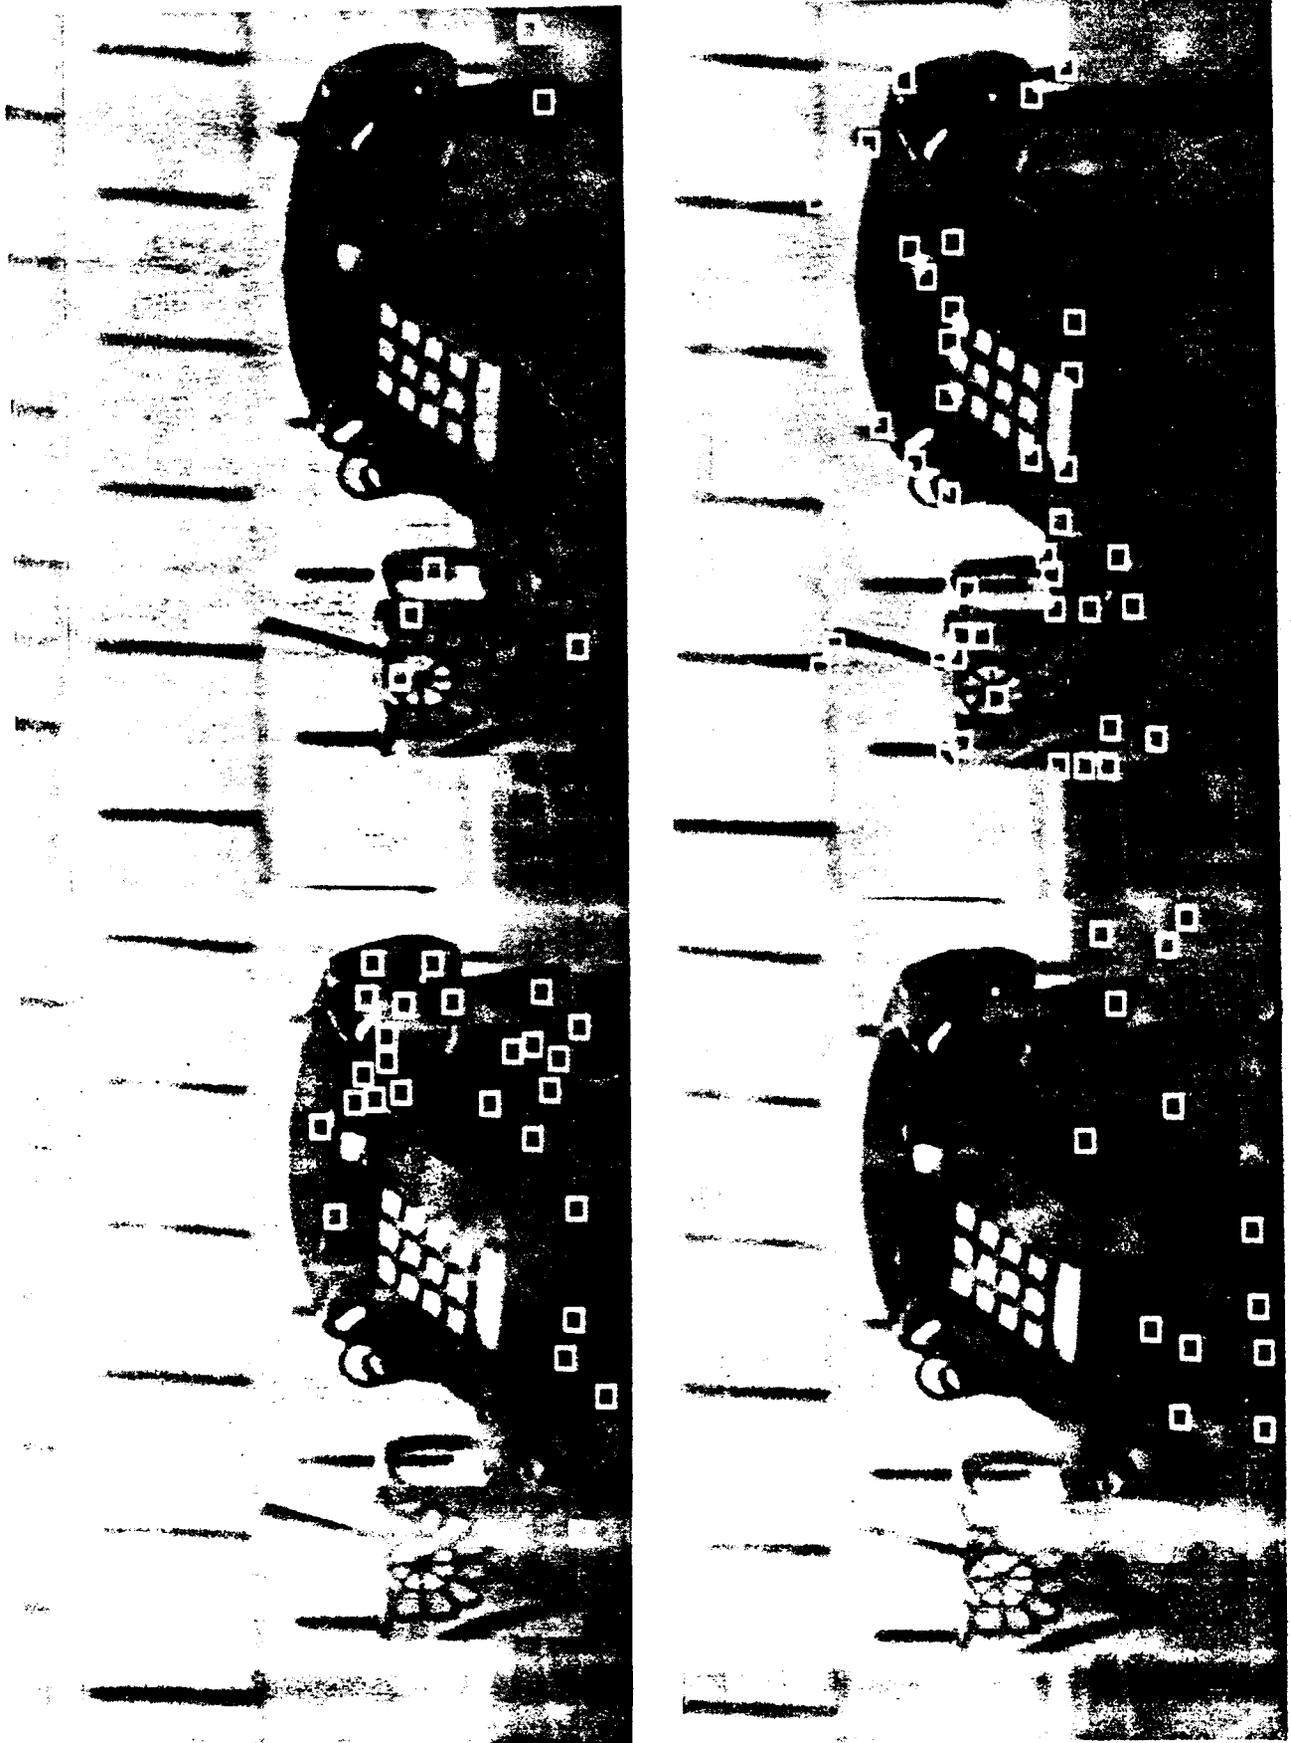


Figure 14: Points picked by Auto-Highest, Auto-Average, Human Expert, and Auto-Weighted (clockwise from top left)