

Digital Stereo Image Matching Techniques

Dr. Marsha Jo Hannah, Senior Computer Scientist
Artificial Intelligence Center, SRI International
333 Ravenswood Avenue, Menlo Park, CA 94025 USA
Commission III/4

Abstract

This paper presents an overview of the various classes of algorithms in use for matching points from one digital image of a stereo pair with the corresponding points in the second image of the pair. These techniques primarily use area-based measures, such as correlation between image patches, or edge-based methods that match linear features in images, but also include the use of feature extractors to match single points in images, as well as global optimization techniques that simultaneously match all points in the two images. This paper also describes an automatic system developed at SRI for stereo compilation; this system uses area-based correlation, but applies this basic technique in a variety of novel ways to develop a disparity model for a given stereo image pair. The techniques used are hierarchical in nature, and incorporate iterative refinement, as well as a best-first strategy, in the matching process. To illustrate these techniques, the results of this system on the Image Matching Test A data set recently distributed by ISPRS's Working Group III/4 are presented.

Introduction

Automatic techniques for the production of three-dimensional (3-D) data via stereo compilation are receiving increased interest for a variety of applications, including cartography [Panton, 1978], autonomous vehicle navigation [Hannah, 1980], and industrial automation [Nishihara and Poggio, 1983]. The first and most difficult step in recovering 3-D information from a pair of stereo images is that of matching points from one digital image of the pair with the corresponding points in the second image. Many computational algorithms have been used in attempts to solve this problem (see [Brady, 1982] or [Barnard and Fischler, 1982] for surveys of the field). These techniques primarily use area-based measures, such as correlation between image patches, or edge-based methods that match linear features in images, but also include the use of feature extractors to match single points in images, as well as global optimization techniques that simultaneously match all points in the two images.

Area matching techniques are the oldest and simplest of the stereo matching algorithms. Each image point to be matched is in fact the center of a small window of points in the first or reference image, and this window is statistically compared with similarly sized windows of points in the second or target image of the stereo pair. The measure of match is either a difference metric that is minimized, such as RMS difference, or more commonly a correlation measure that is maximized, such as mean- and variance-normalized cross-correlation [Hannah, 1974]. Since comparison of a given reference window to every possible target window is computationally expensive, various heuristics have been developed to limit the area that must be searched. In addition to the well-known epipolar constraint, these techniques have included extrapolation from already computed neighboring disparities [Panton, 1978], the use of image hierarchies [Moravec, 1980], and successive iterations of correlation and interpolation [Quam, 1984]. Correlation works well most of the time, but encounters difficulties when the two images are taken from very different viewpoints, of a scene that does not contain adequate visual texture, or of a scene with many depth discontinuities. Even in these instances, and in the presence of image noise, correlation degrades gracefully—it usually continues to find the correct answer, but with reduced confidence measures.

Studies of human vision [Marr and Poggio, 1976] led to the development of edge-based methods, in which linear features are first extracted from the images by an edge operator [Hueckel, 1971; Hildreth, 1980], then matched using the epipolar constraint. Because the processing to extract edges throws away much of the information in the image, many heuristics have been developed to overcome the resulting match ambiguities. These include *a priori* modelling of the scene [Arnold, 1978], multiresolution coarse-to-fine strategies [Grimson, 1981], and longest-first prioritization of the order of edge matching [Baker, 1985]. Because edges are usually somewhat sparse in the image, depths in areas between edges are filled in either by interpolation [Grimson, 1981] or by algorithms that match the image intensities between edges, via dynamic programming techniques [Baker and Binford, 1981; Ohta and Kanade, 1985]. Most edge matching algorithms rely on the relative sparsity of edges, and thus tend to be confounded by images with densely textured areas or moderate levels of image noise, which is precisely where area-based matching excels. For this reason, edge matching should be regarded as complimentary to, rather than as a competitor with, area-based matching. A simple experiment in fusing these two techniques [Baker, 1985] showed vastly improved results over either technique used alone.

Edge-based matching is probably the most-studied form of feature-based matching, but other types of features have also been used for matching. Most of these use a feature operator to select image points that have distinct intensity patterns (corners, centers of small circular areas, etc.), then match the feature vectors for these points, including not only intensities and feature shapes, but also such concepts as the “seldomness” of the features [Förstner, 1986]. Some algorithms find multiple likely matches for each point, then use a relaxation process to disambiguate the results [Barnard and Thompson, 1980]. An extension of feature matching is the detection and matching of regions within images [Price, 1984]. Feature detection, like edge detection, works well in clear, uncluttered images, but suffers when the images are noisy or highly textured.

Most matching algorithms begin by matching individual points in isolation from their neighbors, which leads to errors that must later be detected and rectified using local consistency. A recent promising approach uses simple global optimization techniques that simultaneously match all points in the two images [Barnard, 1987]. The measure of match here is an energy function combining difference in pixel (picture element) intensities between the two images (image similarity) as well as local differences in disparity (scene continuity); the global energy is driven to a minimum at successive levels of the image hierarchies, to limit the search space. Results on aerial imagery are encouraging, although the system does not handle depth discontinuities well, and may be sensitive to image noise.

In implementing our system, we chose to base it on area-based techniques, because of their robustness and wide range of applicability over image types. These attributes were again demonstrated by our results on ISPRS’s test imagery.

Description of SRI’s Stereo System

Over the past five years, SRI has integrated and improved existing pieces of stereo software into a baseline system for automated, area-based stereo compilation. The system operates in several passes over the data, during which it iteratively builds and refines its model of a portion of the 3-D world represented by a pair of images.

The first step in our matching process is to select a set of well-scattered windows in one image, such that each window contains sufficient information to produce a reliable match. To accomplish this, a statistical operator is passed over the image; this is a product of the image variance and the minimum of ratios of directed differences (hence edge strength) over windows of the specified size [Hannah, 1980]. Local peaks in the output of this operator are recorded as the preferred places to attempt the matching process (Figure 1). The motivation behind this operator is that it penalizes windows with low information and windows whose only information

is contained in strongly linear edges, as both of these situations cause difficulties in obtaining the correct match via area-based correlation. The chosen windows are characterized by their center points, which are referred to as “interesting points” [Moravec, 1980]. To ensure that the selected windows are well-scattered in the image, the image is divided into a grid of subimages, and the relative ranks of the interesting points within their grid cell are recorded; this permits the most interesting points in each area to be matched first.

Whether or not point (x_1, y_1) in the first image I_1 is matched by point (x_2, y_2) in the second image I_2 is determined by computing the cross-correlation, normalized by both mean and variance, over windows surrounding the points [Hannah, 1974]. The matching point is taken to be the point in I_2 with highest correlation, as located by one of several search algorithms.

Our system employs several different matching algorithms. The underlying strategy is to begin with a few points that are most likely to be matchable (based on their “interest,” i.e. information content); these are matched by very global, but very conservative, search algorithms. Each successive algorithm operates on less promising points, but uses more information from matches made at previous levels to constrain the search to smaller and smaller portions of the epipolar line, until eventually all interesting points have been processed. All of our matching algorithms use image hierarchies to some extent. Pixels in each reduced image of the hierarchy are produced by convolving the parent image with a Gaussian, then sub-sampling [Burt, 1980]; images are almost always reduced in size by a factor of 2 at each step of the hierarchy (Figure 2).

The first matching algorithm, unconstrained hierarchical matching, assumes that nothing is known about the relative orientations of the images, other than that they cover approximately the same area, at about the same scale, with no major rotation between the images. Each specified point (usually the most interesting point in each grid cell) is matched using an unguided hierarchical matching technique [Moravec, 1980]. This technique begins with a point in the largest image (the 240x240 left image of each of the test sets) and numerically traces that point back through that image’s hierarchy by repeatedly halving the coordinates of the point until it reaches an image that is approximately the size of the correlation window. It then uses a two-dimensional spiral search, followed by a hill-climbing search for the maximum of the correlation between the image windows [Quam, 1971]. This global match is then refined back down the image hierarchy; that is, the disparity at each level (suitably magnified to account for relative image scales) is used as a starting point for a hill-climbing search at the next level (Figure 2). The correlation window size remains constant at all levels of the hierarchy, so the match is effectively performed first over the entire image, then over increasingly local areas of the image. This technique permits the use of the overall image structure to set the context for a match; the gradually increasing detail in the imagery is then followed down through the hierarchy to the final match.

In this matching technique, as in all the others we use, matches must pass fairly strict tests in order to be considered correct, and only the successful matches are recorded for further use. At any level in the hierarchy, matches with poor correlation (compared either against an absolute threshold or with respect to an autocorrelation-based threshold [Hannah, 1974]) are discarded, as are matches that fall outside of the image. Each match must also be confirmed by back-matching; that is, if we have found that point (x_1, y_1) in the first image I_1 is best matched by (x_2, y_2) in the second image I_2 , we then repeat the entire matching algorithm, starting with (x_2, y_2) in I_2 and searching for the point (x'_1, y'_1) in I_1 that best matches (x_2, y_2) . If (x_1, y_1) and (x'_1, y'_1) differ by more than one pixel, the match is discarded as being unreliable. Figure 3 shows the result of applying this technique to the most interesting point in each grid cell.

We next calculate a simplistic relative camera model from the set of point pairs produced by unconstrained hierarchical matching. This is accomplished by searching for five angles that describe the relative positions and orientations of two ideal pinhole cameras [Hannah, 1974]—the azimuth and elevation of the second camera’s focal point with respect to the first camera; and




Figure 1: Results of interest operator on Test 1 (Car I).




Figure 2: Example of unconstrained hierarchical matching.

the pan, tilt, and roll of the second camera's axes with respect to those of the first. The object of the search is to minimize the error between (x_2, y_2) in I_2 and the epipolar line produced when (x_1, y_1) in I_1 is projected into space, then into I_2 through the hypothesized pinhole cameras. The search proceeds by a linearization of the equations and their analytic derivatives [Gennery, 1980]. Once a solution is found, the reliability of each matched point is assessed. Points that appear to contribute too much error to the solution are removed from the calculation, and the solution is redone. This process will reach a successful conclusion if a subset of points is found to converge to a consistent model, or it will report failure if too many points are rejected.

The next technique to be applied is epipolar constrained hierarchical matching. Having determined the camera parameters, we now know the manner in which a point in the first image projects to a line in the second image—the epipolar constraint. This constraint allows us to cut the search from two dimensions (all around the point) to one dimension (back and forth along the epipolar line) at each level of the hierarchy. In all other respects, epipolar constrained hierarchical matching proceeds very much like unconstrained hierarchical matching, with the additional match-evaluation criteria that matches must lie within a specified distance of the epipolar line. This technique is used on any unmatched points among the two most interesting point for each grid cell.

Once a good basis of reliable matches has been found, these matches can be used as “anchor” points for the anchored matching technique, which again uses the grid cells in the image. A given point will lie in some grid cell; the closest matched points should lie in that cell or in one of the eight neighboring cells. Under the assumption that the world is generally continuous, a point would be expected to have a disparity similar to that of its neighbors. Thus, the disparity for a point is expected to lie in the interval of the disparities of the well-matched points in the current and neighboring cells. This disparity interval is used along with the epipolar constraint to perform a very local search for the match to a point, perhaps proceeding one or two levels up the image hierarchy, to provide context for the match. All matches are required to pass the same tests discussed for the hierarchical matching algorithms described previously, including the back-matching test. Figure 4 shows all of the interesting points that were matched by these various techniques.

Our system also can produce matched points on a regularly spaced grid, if desired. This matching algorithm also uses the anchored match technique, searching along specified portions of the epipolar line, to calculate matches for the user-specified grid of points in the first image. (Figure 5). However, holes can result if a grid point does not have suitable information for matching, and again, only matches that pass all the tests are recorded. This highlights a problem with matching a grid of points—not all areas of an image have information suitable for matching, and forcing a match at such areas can lead to poor results, so matching on a grid must be used with caution.

Results on Image Matching Test A

To test our system, we participated in ISPRS Working Group III/4's Test A on Image Matching. For each of the 12 image pairs, we attempted to perform Standard Task B—determination of the parallaxes at selected points—which is what our system does best; for a few of the images, we also attempted to determine the parallaxes at a grid of points. In most cases, the images were run with the standard parameters for the system, which had been tuned to process a high-quality, 1024x1024 aerial image pair. Because of incompatibilities in format, we did not use the camera information given with each test image pair, or any other *a priori* information; we used the raw images, without transforming to normal images, or doing any other resampling.

For the most part, the matching proceeded routinely, using the standard procedures and parameters. One parameter—a threshold that indirectly controls the number of interesting




Figure 3: Results of unconstrained hierarchical matching.




Figure 4: Test 1 (Car I)–Parallaxes at selected points.




Figure 5: Test 1 (Car I)–Parallaxes at a grid of points.

points that the system has to work with—was changed for each data set; this was done to produce between 100 and 300 points per image, as requested for the test. In the paragraphs that follow, we outline processing steps that deviated from the usual and comment on the overall results for each test image pair. In the figures that illustrate our results, matches are shown with one of two different marks; the smaller marks correspond to matches with correlations less than 0.3. The system discarded matches about which it was not confident, but we have done no editing of its final results.

On the images for Test 1 (named Car I), the system produced 335 interesting points, of which it matched 330 (Figure 4). We also attempted to match a grid of points, which produced 507 matches (Figure 5). The interesting points are fairly well scattered throughout the image, providing a good base of matches for further processing. The grid of points is fairly complete. From a brief visual inspection, all of the matches appear to be reasonably correct.

On Test 2 (Quarry), the system produced 275 interesting points, of which it matched 249 (Figure 6). We also attempted to match a grid of points, which produced 437 matches (Figure 7). Again, the interesting points are fairly well scattered throughout the image, providing a good base of matches for further processing, although there are a few blank spots. The grid of matches has more holes than the previous example, but is still fairly complete. All of the matches appear to be reasonably correct. The foreground post was found to be interesting, but those matches were at the disparity of the background, with degraded correlations.

On Test 3 (Olympia I), the system produced 171 interesting points, of which it matched 139 (Figure 8). Because of the sparse information in this image, we did not attempt to produce matches on a grid. Our system had considerable trouble with this image pair, because of the general lack of information other than the ambiguities caused by the many identical buttons. Performing unconstrained hierarchical matching on just the single most interesting point in each grid cell did not produce enough matches to form a camera model, so we started over, asking for the four most interesting points per cell to be attempted. Of the 25 matches produced this way, one was clearly wrong, which caused the camera model solver to be unable to converge to a model, and hence it was unable to edit the point out. In this instance (the only time we did so for any test), we removed the offending point by hand, then tried the camera model solution again, with good results. (If our model solver included the RANSAC technique [Fischler and Bolles, 1980], we believe it could have proceeded without intervention.) To be consistent with the unconstrained matching, the epipolar hierarchical matching was also done for any as-yet unmatched points amongst the four most interesting points per cell, rather than the usual two most interesting points. Doing an anchored match with the normal parameters seemed to miss a lot of points, so we asked for a second pass of anchored matching, using all previous matches as anchors, doubling the number of grid cells in each dimension, and constraining the search to just the highest level of the image hierarchy. This produced a fairly good base of matches, which appear to be reasonably correct.

On Test 4 (South America), the system produced 309 interesting points, of which it matched 169 (Figure 9). We also attempted to match a grid of points, but this resulted in more “holes” than “grid,” so we aborted the effort. The matched points are fairly well scattered throughout the image. Because of the poor image quality, most of the correlations were fairly low, which accounted for most of the points that the system said could not be matched. The better matches that were retained appear to be mostly correct, although some of the matches appear to be off by a little. Because to the fuzziness of the images, it is difficult to assess match correctness.

On Test 5 (Bridge), the system produced 219 interesting points, of which it matched 180 (Figure 10). Based on our difficulties with Test 3, we began by trying to match the 2 most interesting points in each cell instead of just the single most interesting point per cell, as is usual. That worked quite well (i.e. the system probably would have worked correctly with its normal settings), and the system operated routinely from there. The matched points are fairly




Figure 6: Test 2 (Quarry)—Parallaxes at selected points.




Figure 7: Test 2 (Quarry)—Parallaxes at a grid of points.




Figure 8: Test 3 (Olympia I)—Parallaxes at selected points.




Figure 9: Test 4 (South America)–Parallaxes at selected points.




Figure 10: Test 5 (Bridge)–Parallaxes at selected points.




Figure 11: Test 6 (Tree)–Parallaxes at selected points.

well scattered along the “wall” of the bridge, and are strung along the lines on the “floor” of the bridge. The matches appear to be reasonably correct, although a couple of the matches appear to be off by one wire in the fence, or by one grid element in the floor, because of the ambiguities of repetitive patterns.

On Test 6 (Tree), the system produced 158 interesting points, of which it matched 109 (Figure 11). The interesting points are strung out along the various branches, marking some but not all of the branch forks and crooks. A few “interesting points” are off to the side of a branch, rather than centered on it; this is a known weakness of our interest operator. Despite the unusual nature of the scene, it was processed with the usual parameters and sequencing of routines. Most of the matches are locally plausible, although some of them have tracked “false intersections”—places where one branch passes in front of another, rather than an actual real-world point.

On Test 7 (Island), the system produced 311 interesting points, of which it matched 260 (Figure 12). The interesting points are fairly well scattered throughout the image, providing a good base of matches for further processing. Most of the matches appear to be reasonably correct, although there are a few questionable ones near the edges of the image.

On Test 8 (Switzerland), the system produced 279 interesting points, of which it matched 196 (Figure 13). The interesting points are fairly well scattered throughout the image. The matches appear to be substantially correct, although they represent a mixture of tree tops and ground points, since our system is unable to distinguish between the two. Many of the correlations are low, possibly because the two images appear to have been recorded at significantly different times of day.

On Test 9 (Car II), the system produced 271 interesting points, of which it matched 168, using the usual processing routine and parameters, despite the wide range of disparities (Figure 14). The matched points are concentrated in a few clusters on high-intensity areas of the image. It is difficult to fuse these images long enough check the results, but most of the matches appear to be approximately correct, with the possible exception of a few matches near the headlight.

On Test 10 (Wall), the system produced 224 interesting points, of which it matched 194 (Figure 15). The matched points are fairly well scattered throughout the image, avoiding blank spots and the very linear edges between faces. The matches appear to be reasonably correct, with a few possible problems near the edge of the image.

On Test 11 (Olympia II), the system produced 197 interesting points (Figure 16). After the difficulty we had with Test 3 (Olympia I), we tried matching the 4 most interesting points in each cell. Of these 133 points, only 4 resulted in matches that our system thought were reliable, based on its local criteria, and none of these matches were actually correct. We therefore abandoned processing on this image pair. We believe that the combination of repetitive structures, the rather different points of view, and the transparency of the dome caused our hierarchical matching techniques to fail. If we had elected to use the accompanying camera information, we might have been able to do some matching, although the transparency would undoubtedly have lead to numerous problems.

On Test 12 (House), the system produced 119 interesting points, of which it matched 92 (Figure 17). The interesting points are poorly distributed throughout the image, avoiding the featureless areas of the driveway, lawn, and roof, and the linear edges of the shadow. The matches appear to be reasonably correct, even in the difficult area where the back of the house falls off to the ground.

Summary

In this paper, we have discussed the various classes of algorithms in use for matching points between the digital images of a stereo pair, including area-based measures such as correlation




Figure 12: Test 7 (Island)–Parallaxes at selected points.




Figure 13: Test 8 (Switzerland)–Parallaxes at selected points.




Figure 14: Test 9 (Car II)–Parallaxes at selected points.




Figure 15: Test 10 (Wall)–Parallaxes at selected points.




Figure 16: Test 11 (Olympia II)–Unable to determine parallaxes.




Figure 17: Test 12 (House)–Parallaxes at selected points.

between image patches, and edge-based methods that match linear features in images, as well as the use of feature extractors to match single points in images, and global optimization techniques that simultaneously match all points in the two images. We have described our automatic system for stereo compilation, using area-based correlation, which applies this basic technique in a variety of novel ways. Our techniques are hierarchical in nature, and use iterative refinement, as well as a best-first strategy, in the matching process, as well as the constraint of backmatching to verify matches. Finally, we have illustrated our techniques by presenting some of our results on the Image Matching Test A data set recently distributed by ISPRS's Working Group III/4.

Overall, we are very pleased with the results of our system on the test images. For 10 of the 12 tests, we were able to handle the image pairs without substantially altering the default parameters or processing sequence. Our only problems came on the images of the Olympia dome; this is not surprising, since our algorithms were designed for use on highly textured natural terrain, not the bland faces of cultural objects. For the most part, our results appear to be reasonably correct, even in the face of large disparity ranges within small areas of the image. As of this writing, we have not received the results of the committee's detailed analysis, which should be most interesting.

Acknowledgements

The research reported herein was supported by the Defense Advanced Research Projects Agency under Contracts MDA903-83-C-0027 and MDA903-86-C-0084, which were monitored by the U.S. Army Engineer Topographic Laboratory. The views and conclusions contained in this paper are those of the author and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or of the United States Government. I would like to thank Harlyn Baker, Lynn Quam, Grahame Smith, and Martin Fischler for their support on this project.

References

- Arnold, R. D., 1978. "Local Context in Matching Edges for Stereo Vision," *Proceedings: Image Understanding Workshop*, Cambridge, MA, May, 1978, pp. 65-72.
- Baker, H. H., 1985. "Reimplementation of the Stanford Stereo System and Integration Experiments with the SRI Baseline Stereo System," unpublished presentation at *Image Understanding Workshop*, Miami Beach, FL, December, 1985; also available as SRI International Artificial Intelligence Center Technical Note 431, January, 1988.
- Baker, H. H., and Binford, T. O., 1981. "Depth from Edge and Intensity Based Stereo," *Seventh International Joint Conference on Artificial Intelligence*, Vancouver, B.C., August, 1981, pp. 631-636.
- Barnard, S. T., 1987. "Stereo Matching by Hierarchical Microcanonical Annealing," *Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy, August, 1987, pp. 832-835.
- Barnard, S. T., and Fischler, M. A., 1982. "Computational Stereo," *ACM Computing Surveys*, Vol. 14, No. 4 (December, 1982), pp. 553-572.
- Barnard, S. T., and Thompson, W. B., 1980. "Disparity Analysis of Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-2, No. 4 (July, 1980), pp. 333-340.
- Brady, M., 1982. "Computational Approaches to Image Understanding," *ACM Computing Surveys*, Vol. 14, No. 1 (March, 1982), pp. 3-71.
- Burt, P. J., 1980. "Fast, Hierarchical Correlations with Gaussian-like Kernels," University of Maryland Computer Science Center Report TR-860, January, 1980.

- Fischler, M. A. and Bolles, R. C., 1980. "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography," *Proceedings: Image Understanding Workshop*, College Park, MD, April, 1980, pp. 71-88.
- Förstner, W., 1986. "A Feature Based Correspondence Algorithm for Image Matching," *International Archives of Photogrammetry*, Vol. 26-III, Rovaniemi, Finland, 1986.
- Gennery, D. B., 1980. "Modelling the Environment of an Exploring Vehicle by Means of Stereo Vision," Ph.D. Thesis, Stanford University, Computer Science Department Report STAN-CS-80-805, June, 1980.
- Grimson, W. E. L., 1981. *From Images to Surfaces: A Computational Study of the Human Early Visual System*, M.I.T. Press, Cambridge, MA, 1981.
- Hannah, M. J., 1974. "Computer Matching of Areas in Stereo Images," Ph.D. Thesis, Stanford University, Computer Science Department Report STAN-CS-74-438, July, 1974.
- Hannah, M. J., 1980. "Bootstrap Stereo," *Proceedings: Image Understanding Workshop*, College Park, MD, April, 1980, pp. 201-208.
- Hannah, M. J., 1985. "SRI's Baseline Stereo System," *Proceedings: Image Understanding Workshop*, Miami, FL, December, 1985, pp. 149-155.
- Hildreth, E. C., 1980. "Implementation of a Theory of Edge Detection," M.Sc. Thesis, Massachusetts Institute of Technology, Artificial Intelligence Laboratory Technical Report AI-579, 1980.
- Hueckel, M., 1971. "An Operator Which Locates Edges in Digital Pictures," *Journal of the Association for Computing Machinery*, Vol. 18, pp. 113-125.
- Marr, D., and Poggio, T., 1976. "Cooperative Computation of Stereo Disparity," *Science*, Vol. 194 (1976), pp. 283-287.
- Moravec, H. P., 1980. "Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover," Ph.D. Thesis, Stanford University, Computer Science Department Report STAN-CS-80-813, September, 1980.
- Nishihara, H. K., and Poggio, T., 1983. "Stereo Vision for Robotics," *Proceedings of the International Symposium of Robotics Research*, Bretton Woods, NH, September, 1983.
- Ohta, Y., and Kanade, T., 1985. "Stereo by Intra- and Inter-Scanline Search Using Dynamic Programming," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 2 (March, 1985), pp. 139-154.
- Panton, D. J., 1978. "A Flexible Approach to Digital Stereo Mapping," *Photogrammetric Engineering and Remote Sensing*, Vol. 44, No. 12, pp. 1499-1512.
- Price, K. E., 1984. "Matching Closed Contours," *Proceedings: Image Understanding Workshop*, New Orleans, LA, October, 1984, pp. 169-175.
- Quam, L. H., 1971. "Computer Comparison of Pictures," Ph.D. Thesis, Stanford University, Computer Science Department Report STAN-CS-71-219, May, 1971.
- Quam, L. H., 1984. "Hierarchical Warp Stereo," *Proceedings: Image Understanding Workshop*, New Orleans, LA, October, 1984, pp. 149-156.
- Smith, G. B., 1984. "A Fast Surface Interpolation Technique," SRI International Artificial Intelligence Center Technical Note 333, August, 1984.