Stereo Error Detection, Correction, and Evaluation

RAKESH MOHAN, MEMBER, IEEE, GERARD MEDIONI, MEMBER, IEEE,
AND RAMAKANT NEVATIA, SENIOR MEMBER, IEEE

Abstract—Even with an increased sophistication of stereo matching techniques, only incremental gains have been made in their performance. The purpose of this paper is to present an algorithm for error detection and correction of disparity, as a process separate from stereo matching, with the contention that matching is not necessarily the best place to exploit all the physical constraints characteristic to stereopsis. As a result of the bias in stereo research towards matching, vision tasks like surface interpolation and object modeling have to accept erroneous data from the stereo matchers without the benefit of any intervening stage of error correction. We present an algorithm which identifies all errors in disparity data which can be detected on the basis of figural continuity and corrects them. The algorithm can be used as a post processor to any edgel based stereo matching algorithm[1] and can additionally be employed to automatically provide quantitative evaluations of the performance of matching algorithms of this class.

Index Terms—**Linear features, matching, stereo.**

I. Introduction

THE relative displacement or *disparity* in the position of objects, as viewed by a pair of eyes, is an important source of depth information for humans. This phenomenon, termed *binocular stereopsis*, can be used to recover some of the 3-D information present in the scene. A survey of methods can be found in [1].

The most difficult task in stereo is that of identifying corresponding locations in the two images. Understandably therefore, research in stereopsis has been focused primarily on the development of matching criteria and matching algorithms. However, the comparable performance of the wide variety of current stereo matching algorithms indicates that use of some of the sources of information in processes other than matching (rather than attempts to incorporate diverse sources of information into a matching criterion), might result in substantial improvement in the overall performance of stereo vision systems. There are physical properties of the stereopsis process which are more amenable to treatment at the level of match disparities. For instance, our utilization of the figural continuity constraint on the disparity data is more precise and simpler than previous attempts at its incorporation in the matching algorithm.

A study of the physics of stereopsis reveals various properties which can be employed in implementations of stereo vision. Of these, the *epipolar geometry constraint*, which restricts the possible matches of a feature to those lying in its epipolar line, has been used effectively to reduce the search space for the matcher. The constraint imposed by the smooth nature of physical *surfaces*[2] has been used by stereo matchers in various guises, for example as principles of depth being continuous almost everywhere [2], [3] and of minimum differential disparity [4]. The working of these algorithms depends significantly on an abundance of surface markings. Their performance quality suffers markedly at occluding boundaries [5], where large changes in disparity can occur, and it is these contours that play an important role in scene understanding.

Intensity discontinuities in images correspond to physical features in the imaged scene. Thus edge contours[3] correspond to surface boundaries or surface markings. These feature contours are connected, therefore the change in depth along them is continuous. As disparity is inversely proportional to depth, we have the following *figural continuity constraint* on disparities:

> Disparity along an edge contour changes smoothly, i.e., there should be no disparity discontinuities along a contour.

The figural continuity constraint was first exploited by Mayhew and Frisby [6], [7]. To apply this constraint, we have to know the functional relationship between the disparity of an edgel and its position along the contour, and for edgel based stereo matching, utilize this information in a consistent fashion across epipolar lines. Stereo matchers incorporating this constraint have been plagued by various deficiencies. A precise formulation of disparity change along contours is missing—it often being unrealistically assumed that all edge contours are nearly parallel to the imaging plane so the disparity along them is about constant [8]-[10]. Algorithms that are able to detect the matches violating figural continuity either reject them [9] or try to choose suitable replacements from alternate matches [8]. The latter is insufficient when the correct

Manuscript received January 16, 1987; revised June 1, 1988. Recommended for acceptance by O. Faugeras. This work was supported by the Defense Advanced Research Projects Agency under Contract F33615-84-K-1404, monitored by Wright Patterson Aeronautical Labs.

The authors are with the Institute for Robotics and Intelligent Systems (IRIS), Departments of Electrical Engineering and Computer Science, University of Southern California, Los Angeles, CA 90089.

IEEE Log Number 8824800.

[1]An *edgel* is an edge having an extent of one pixel. *Edgel based stereo* uses edgels as the features to match.

[2]Surface here is a closed area containing no orientation or depth discontinuities.

[3]In this paper we call a contour a segment of a curve bounded by corners or the ends of the curve.

match is missed during the selection of possible matches or is missing from the matched image. Attempts at passing the continuity information between epipolar lines have been faulty [10] or extremely expensive [11]. Further, the matching algorithm may have to resolve conflicts between matches preferred by figural continuity to those postulated by the similarity of the feature elements.

If we turn our attention from matching to the resulting disparities, most of these problems vanish. We consider edgel based stereo and classify matching errors into two types on the basis of their detectability on the basis of figural continuity alone (Section II). Then we prove that disparity varies linearly along *linear segments*[4] (Subsection III-A). An algorithm is presented which, given a linear segment, first rejects the disparities violating figural continuity and then computes the linear function describing the disparity change along that segment (Subsection III-C). This function is then used to correct disparities at the erroneous matches and fill in missed matches. By divorcing the application of the figural continuity constraint from the matching phase, we avoid the last two problems mentioned above. This also allows the use of this algorithm as a postprocessor to any edgel based stereo system, independent of the particular matching algorithm used. As our algorithm is applied to disparity data, it does not necessitate any changes to the stereo matching system. Additionally, our algorithm is very efficient—it is linear in the number of edgels in the image.

Since we use figural continuity along segments as our basis for error correction and detection, our algorithm is not applicable, in general, the stereo algorithms using linear segments as the basic feature for matching [4]. In the segment-based stereo matching, the use of figural continuity is *implicit* in the calculation of disparity along a matched segment. However, existing algorithms for matching segments are very complex [4]. Our algorithm has the benefit of allowing any edgel-based stereo system to exploit linear segments in a simple yet effective way.

We show our algorithm to be useful in another problem area in stereo—which might appear unrelated at first—that of the *evaluation of matching algorithms*. Usually authors provide a quantitative evaluation of their stereo algorithms by manually counting the number of matching errors their algorithms commit. This task can be automated by modifying our scheme such that it keeps a count of the errors it detects. This scheme of quantitative evaluation is demonstrated by a comparison of four matching algorithms on its basis (Section IV).

II. Errors in Line-by-Line Stereo

It is our observation that in any stereo matching method, a lot of edgels are assigned wrong matches. Wrong matches could be incorrect null matches (i.e., edgels assigned no match), or matches to a wrong edgel. Edgels might not be matched because the match evaluation function used by the stereo matcher assumes that the edgel is spurious, or that its matching edgel in the other image is missing or that the edgel is occluded in the other image. Wrong edgel matches are of the following two types:

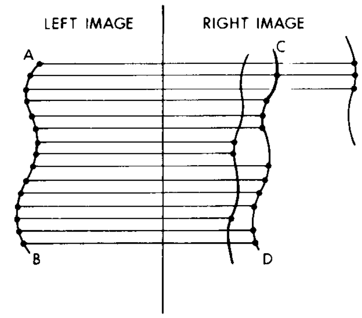

Fig. 1. Type I (local) errors.

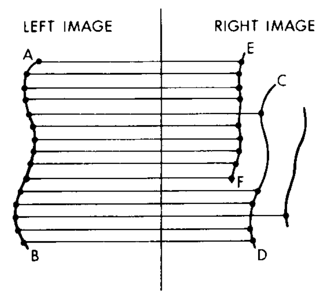

Fig. 2. Type II (global) errors.

- *Type I (local) Errors:* In Fig. 1 the edgel pairs matched by a stereo algorithm are shown linked by the epipolar line. The contours represent the segments detected in the image. The figure shows that more edgels of the segment *AB* are assigned matches to the correct segment, *CD*, than to any single wrong segment. These wrong matches could arise due to the fact that information along a single row of pixels may not be sufficient basis for matching and so there will be some wrong matches on an epipolar line. Also the real epipolar line could be locally distorted due to the imaging device and conditions [5] or due to noise.

Type I errors can be detected and corrected on the basis of the continuity constraint.

- *Type II (global) Errors:* Fig. 2 shows another type of wrong match. Here more edgels of the segment *AB* are assigned matches to a wrong segment *EF* than to the correct segment *CD*. We term such errors as Type II errors. These errors reflect that the function for evaluating the quality of match used by the stereo algorithm prefers the wrong segment as a better match. This is a drawback of the match evaluation function. We do not believe that any given evaluation function can always avoid Type II errors.

Type II errors can not be detected (or corrected) on the basis of the continuity constraint. In fact, all stereo algorithms which use figural continuity along segments, including segment based stereo, can directly deal with only Type I errors.

[4]Linear segments are piecewise linear representations of edge contours.

This classification of errors holds for all types of edge contours. However, from this point on, we will be dealing exclusively with linear segments.

III. Detecting and Correcting Type I Errors

A. Disparity Change Across Linear Segments

Let us now consider the special case of linear segments. A linear segment in a projection of a scene is the image of a linear feature (or a linear approximation of a feature) in the 3-D scene, ruling out accidental alignments. In a stereo image pair, even such accidental alignments in one image will be revealed in the other image (in our current implementations, we do not detect any accidental alignments).

Consider Fig. 3 which shows a linear segment AB in the left image and the corresponding segment CD in the right image. $C'D'$ is the position of CD when the two images are registered such that edgels with zero disparity coincide. The disparity (in the left image) at A is d_A and B is $-d_B$ (associating positive disparity with positive height above the ground).

From

$$\Delta AC'E = \Delta EBD' \tag{1}$$

we have

$$\frac{d_A}{|AE|} = \frac{d_B}{|EB|} = \sin\theta. \tag{2}$$

It is obvious from (2) that disparity varies in a linear proportion to the length of a linear segment (the proportionality constant changing with the orientation of the segment in 3-D). We utilize this constraint for linear segments to detect and correct Type I errors.

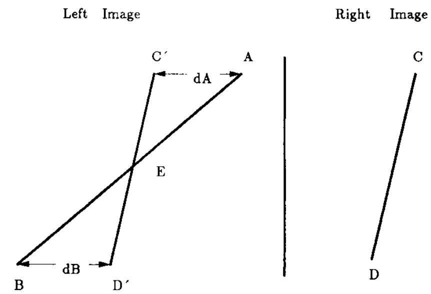

Fig. 3. Matching linear segments in alstereo pair.

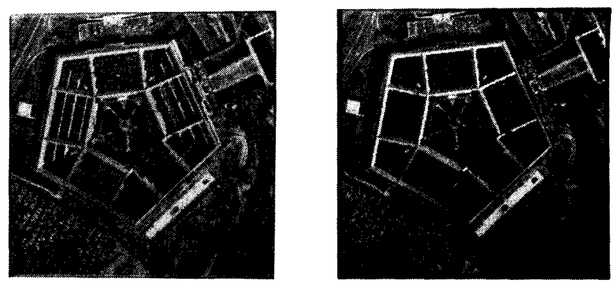

Fig. 4. Aerial image of the pentagon.

B. Obtaining the Disparity Map

The stereo error detection and correction is demonstrated on a stereo pair of the Pentagon as shown in Fig. 4. The edges are obtained using LINEAR, a linear feature extraction system developed by Nevatia and Babu [12]. The disparity at each edgel point is obtained by applying the algorithm proposed by Ohta and Kanade [11] for intrascanline matching. Note that edgels and intensity values are required by this algorithm. For display purposes, the disparities so obtained are interpolated linearly along epipolar lines and the resulting dense disparity map is displayed in Fig. 5 rendered as an intensity array with areas further away from the viewer appearing lighter. The *streaks* in the disparity map are indicators of *matching errors*—in a perspective plot these appear as spikes and dips. For our purpose, any edge detection technique and edgel based stereo algorithm could have been used to obtain the disparity map.

C. Detection and Correction of Errors

LINEAR is then used to fit linear segments to the all the detected edges and the resulting segments are shown in Fig. 6. First a linear segment is picked. We can associate each linear segment with the edgels it represents. Working in a length–disparity (l-d axes) coordinate frame, the disparity for each matched edgel belonging to this segment is plotted as a function of the distance of the edgel from one end of the segment (the length varies from zero to the full length of the segment along the l axis, and the range of disparities along the d axis corresponds to the search range employed by the stereo algorithm). Since disparity varies linearly along a linear segment, if all edgels are matched correctly, all the points plotted should fall on one straight line, as shown in Fig. 7.

Fig. 8 is an actual plot of a disparities at the edgels of a typical segment from Fig. 6. If we can interpolate a line through this data which would correspond to the correct disparities, we could not only *detect* and *correct* wrong matches but also *fill in* disparity information at edgels which were not matched and at edgels belonging to occluded portions of the segment.

We have to throw away the points plotted due to wrong matches before we can interpolate a line through the disparity values. If more edgels are matched to the correct segment than to any other single wrong segment (Type I error) then we can detect the errors in the following way. We fit thin strips of all orientations and locations to the disparity versus length data plotted. From the definition of Type I errors, it follows that the strip which has the maximum number of points in it has the disparity values corresponding to the correct matches, the rest are erroneous. The reason for choosing thin strips instead of straight lines is that due to the discrete location of the edgels and the error in the location of edgels, all the points from correct matches do not fall exactly on a line.

Fig. 5. Original disparity map obtained.

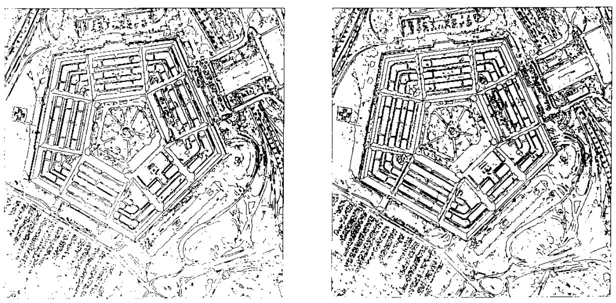

Fig. 6. Linear segments detected in Fig. 4.

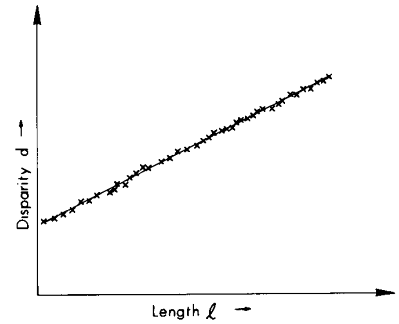

Fig. 7. Plot of ideal disparity versus length on a segment.

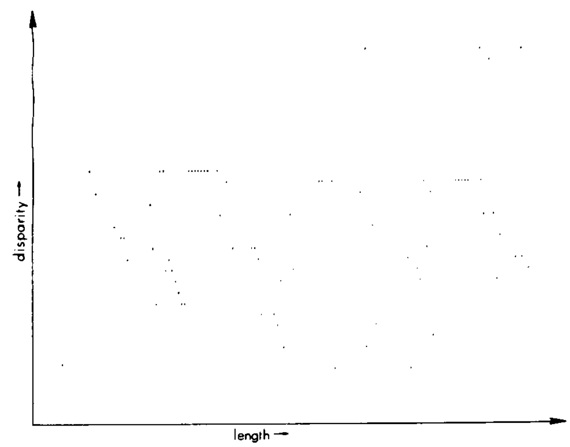

Fig. 8. Plot of real disparity versus length on a segment.

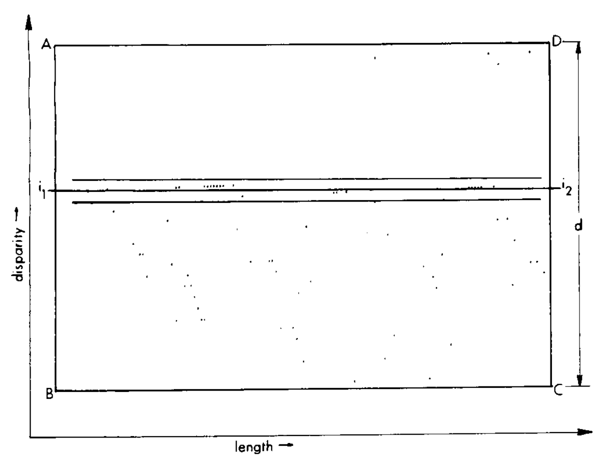

Fig. 9. Thin strip containing the correct disparities.

To fit the thin strips to the data we use a version of the modified Hough transform technique proposed by Wallace [13]. Consider Fig. 9: the window *ABCD* of the *d-l* plane contains all the plotted disparities. The thin strip fitted on the correct data points has to intercept *AB* and *CD*. We choose the intercepts i_1 and i_2 made by the line running through the center of the strip at *AB* and *CD* as the parameters for the Hough transform. The width of the strip will depend on the algorithm used for edge detection. If the edge detector cannot resolve between two edges less than p pixels apart, then the width of the strip has to be less than $2p$. We use a strip three pixels wide and we allow an overlap of one pixel between two adjacent parallel strips. We will then have to consider the intercepts only at every other pixel along *AB* and *CD*. Say the range of disparity in the image is d. Then the search space for the Hough transform is $d^2/4$. Although Fig. 9 shows the correction process for a segment which is horizontal, i.e., parallel to the image plane, our technique is neither limited to nor biased towards horizontal segments and works equally well for nonhorizontal segments.

Once the correct matches have been identified as the points contained in the strip with the maximum number of points, we fit a line through their disparities which minimizes the squared error. From this line we can get the disparities of all the edges on the segment. These new disparities are stored as real numbers since we can get subpixel disparity values. We also store the disparities at the begin and the end points of segments for display purposes.

This process is repeated for all linear segments in the image. The error detection phase reported 24.37 percent of the matches as being erroneous. Fig. 10 shows the gray scale rendering of the corrected disparities after linear interpolation. No smoothing or filtering is used. The structure in the top right corner is a bridge over the road. The ramps at the left corner of the pentagon and the image

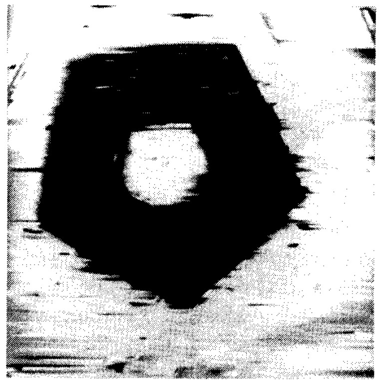

Fig. 10. Corrected disparity map.

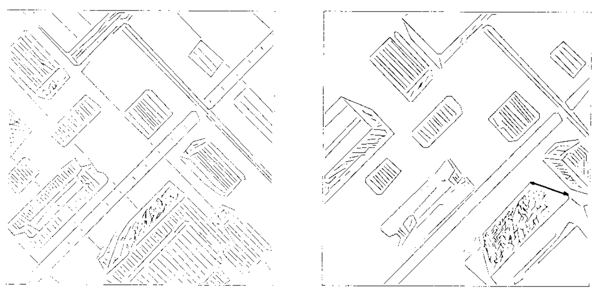

Fig. 11. Linear segments detected in CDC image.

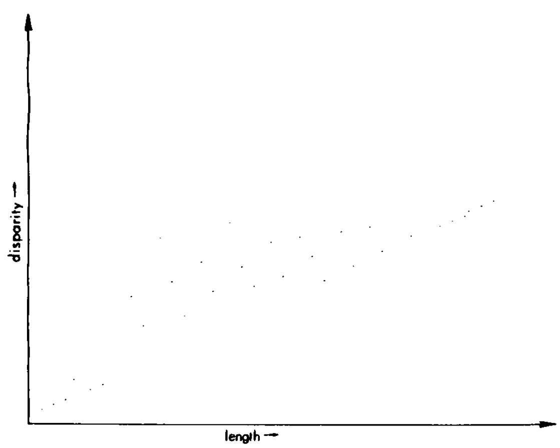

Fig. 12. Plot of disparity versus length of segment marked in Fig. 11.

boundaries are remnants of the linear interpolation of the disparity map. A few streaks visible correspond to Type II errors and almost all of these can be attributed to very short segments.

The next example shows that there can be large disparity variations along important edge contours. In Fig. 11, disparities along the marked segment vary over almost the complete range of disparities for the stereo pair. Fig. 12 shows the errors committed during matching. Note that

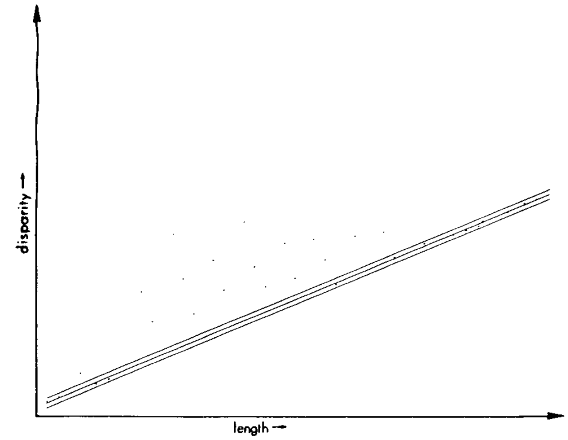

Fig. 13. Correction of disparities on segment marked in Fig. 11.

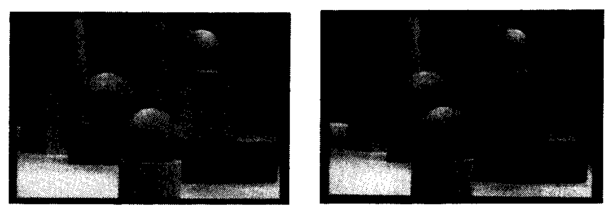

Fig. 14. An indoor scene.

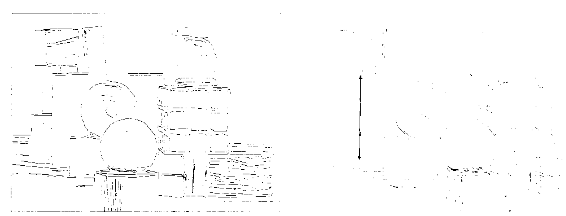

Fig. 15. Linear segments detected in Fig. 14.

the errors are systematic rather than random. Fig. 13 shows the correction performed.

Next, a stereo image pair from a different domain, indoor scenes, is considered. Figs. 14 and 15 show the intensity images and the edges detected, respectively. Due to the sparsity of the edges and the large, sharp changes in disparity over the scene, simple linear interpolation does not give a meaningful display of the disparity data. We, therefore, choose to display in Fig. 16 an orthogonal view of the corrected linear segments in 3-D space.

Twenty-five percent of the matches were in error as reported by the error detector. A typical segment (emphasised in Fig. 15) is chosen to display the errors in the match and their correction in Fig. 17.

At this stage, we are not able to detect Type II errors. Information from other sources which provide depth information, such as monocular cues, shape from shading, smoothness constraints on surfaces, etc. might be useful towards detecting Type II errors.

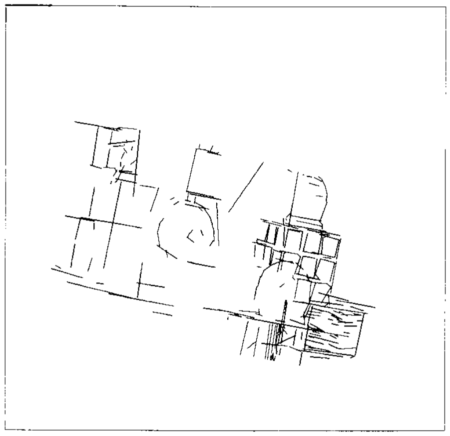

Fig. 16. Orthogonal view of the corrected disparities.

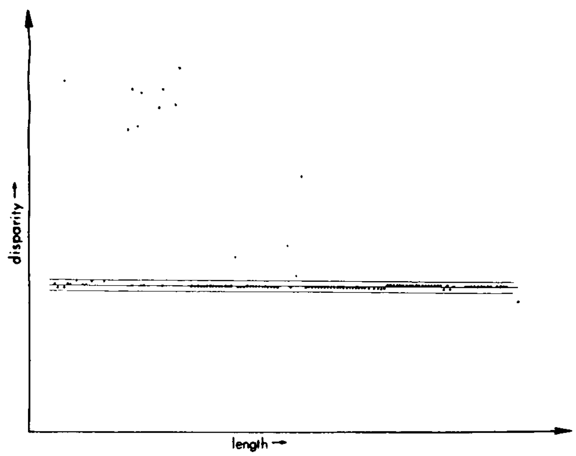

Fig. 17. Correction of disparities at the linear segment marked in Fig. 15(b).

IV. Quantitative Comparison of Matching Algorithms

A reliable quantitative measure of the performance of stereo algorithms would prove useful in evaluating stereo algorithms, improving match evaluation functions and choosing stereo algorithms for special (restricted) domains.

An obvious performance measure is the percentage of matching errors and the percentage of edgels matched. The percentage of edgels matched can be easily obtained from the stereo algorithm itself. Used alone it could be misleading as it would not be known how many of those matches were incorrect. To calculate the percentage of errors, we need an algorithm to consistently identify the wrong matches. As we have seen in Section III-C, we can locate Type I errors consistently.

The algorithm to calculate the percentage error is:

1) Link the edgels into linear segments.
2) Plot the disparities obtained along the length of the segment.
3) Fit thin strips to them to identify the correct matches.
4) Accumulate the number of correct and incorrect matches.
5) Repeat steps 2)–4) for each segment.
6) Percentage error = 100* wrong-matches/total-matches.

We take a typical edgel based stereo algorithm using dynamic programming to compute the best matched edgel sequence along an epipolar line, i.e., the sequence which has the minimum total cost of matching edgel pairs [11]. Four different match evaluation functions, i.e., functions for evaluating the cost of a match, are compared. These functions (except the first one) do not correspond to any standard stereo matching algorithms and have been selected only to demonstrate the evaluation algorithm.

Tables I and II compare the performance of the dynamic programming based matching algorithm using four different match evaluation functions on two stereo image pairs.

The following are among the measures used for the evaluation:

- *Total number of segments selected:* We use segments larger than a fixed threshold to perform error detection.
- *Total number of segments processed:* From the selected segments, only segments having more than a fixed number of matched edgels, and correct matches, are used for the error statistics (and for correction) to ensure reliability.
- *Percentage matched edgels processed:* This is a percentage ratio of the matched edgels which were on the processed segments to all the edgels matched.
- *Percentage error:* This is the percentage of matched edgels which were in error among the matched points processed.
- *Percentage edgels corrected:* Number of edgels which were either corrected or had their disparities filled in from all the edgels on the processed segments.

The match evaluation functions compared were:

- *Cost Function I:* This is the cost function used by Ohta and Kanade [11].

$$m = \frac{1}{2} \times \left(\frac{1}{k} \sum_{i=1}^{k} a_i + \frac{1}{l} \sum_{j=1}^{l} b_j \right) \tag{3}$$

where $a_1 \cdots a_k$ and $b_1 \cdots b_l$ are the intensity values at the pixels on the left of the matched edgels.

$$\sigma^2 = \frac{1}{2} \times \left(\frac{1}{k} \sum_{i=1}^{k} (a_i - m)^2 + \frac{1}{l} \sum_{j=1}^{l} (b_j - m)^2 \right) \tag{4}$$

$$\text{cost} = \sigma^2 \times (k^2 + l^2)^{1/2} \tag{5}$$

- *Cost Function II:* To function I we add the restriction that matches are considered only among edgels whose direction do not differ from each other by more than 30.0 degrees. The directions used are the orientation of the segment the edgel belongs to, and not local edge orientations.

TABLE I
PERFORMANCE OF COST FUNCTIONS ON IMAGE I

	COST FUNCTION I	COST FUNCTION II	COST FUNCTION III	COST FUNCTION IV
% edges matched	88.83	88.38	59.07	78.89
% segments selected	38.45	38.45	38.45	38.45
% segments processed	24.14	23.85	18.33	14.31
% matched edge points processed	83.85	78.29	68.90	53.07
% error	3.85	3.83	8.04	31.92
% edge points corrected	14.91	15.13	40.61	45.82
maximum % error for a segment	30.00	27.27	50.00	61.64

TABLE II
PERFORMANCE OF COST FUNCTIONS ON IMAGE II

	COST FUNCTION I	COST FUNCTION II	COST FUNCTION III	COST FUNCTION IV
% edges matched	77.52	74.91	56.95	70.53
% segments selected	37.31	37.31	37.31	37.31
% segments processed	20.15	20.15	14.18	11.94
% matched edge points processed	70.38	70.77	54.16	33.69
% error	8.01	7.52	17.01	20.45
% edge points corrected	14.89	16.11	41.11	42.95
maximum % error for a segment	46.88	44.44	58.82	52.17

- *Cost Function III:* This function is formulated to favor matches among edgels with similar orientations and with similar interval lengths on their left between them and their preceding match.

$$\text{cost} = \left(\exp\left(\left|\theta_l - \theta_r\right|/c_1\right) \times c_2 + \left(\left|\alpha_l - \alpha_r\right| \div (\alpha_1 + \alpha_2)^2 \times c_3\right)\right) \times (\alpha_l + \alpha_r). \qquad (6)$$

Where θ are the edge orientations, α the interval lengths, and c_i constants.

- *Cost Function IV:* For this function we use the intensity intervals lying on the left of the two edge points. The cost of matching the two intervals is the best cost of matching them using an area based stereo algorithm. The area based stereo algorithm used is one proposed by Leguilloux [14] (details of this algorithm are not relevant to our discussion).

Tables I and II list the performances of four algorithms on two stereo pairs of scenes of varying complexity. The performance measures for the matching algorithms, on these scenes, agree on the fact that the performances of Cost Functions I and II are comparable, while those of Cost Functions III and IV are much worse. A wide variety of scenes and algorithms need to be considered before we can make claims to the stability of the performances measures of this evaluation algorithm, but the comparable figures for the two different scenes is an indication of it.

V. CONCLUSION

We have presented an algorithm, which working independently of the stereo matching, improves the accuracy of the disparity data. Our observation about disparities along linear segments gives us a precise formulation of the figural continuity constraint, which had been missing from previous implementations [15], but the key gains come from the fact that rather than burdening the matching algorithm with the task combining epipolar geometry, similarity of matched features and figural continuity, we choose to work with disparity data directly. This is also an indication that the match disparities could be a suitable domain for treating the problem of Type II errors.

The time complexity of the proposed algorithm is *linear* in the number of edgels processed. Therefore, it is a very efficient way of ensuring figural continuity along segments. Although the algorithm does not use the continuity constraint during the matching process, we believe that since a substantial number of the edgels can be processed by this algorithm, we can afford to ignore the other matched edgels. Even if the corrected disparities are used as fixed constraints for a second pass over the image pair, the algorithm would both be computationally cheaper and would use more continuity constraints than the one described in [11].

We have been able to demonstrate that we can process a large number of the matched edgels and detect Type I errors in them. We can also correct these errors and fill in disparities for edgels not matched. However, this algorithm is weak in the following areas:

- Not all matched edgels are processed. This is due to the fact that at each segment we need some minimum number (three or four is sufficient) of correct matches before we can confidently interpolate disparity for the whole segment.
- The problem of Type II errors has not been addressed.

REFERENCES

[1] S. Barnard and M. Fischler, "Computational stereo," *ACM Comput. Surveys*, vol. 14, no. 4, pp. 553–572, Dec. 1982.
[2] D. Marr, *Vision.* San Francisco, CA: Freeman, 1982.
[3] W. E. L. Grimson, *From Images to Surfaces.* Cambridge, MA: M.I.T. Press, 1981.

[4] G. Medioni and R. Nevatia, "Segment-based stereo matching," *Comput. Vision, Graphics, Image Processing*, vol. 31, pp. 2-18, 1985.
[5] W. E. L. Grimson, "Computational experiments with a feature based stereo algorithm," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 7, no. 1, 1985.
[6] J. E. W. Mayhew and J. P. Frisby, "The computation of binocular edges," *Perception*, vol. 9, pp. 69-86, 1980.
[7] —, "Psychophysical and computational studies towards a theory of human stereopsis," *Artificial Intell.*, vol. 17, 1981.
[8] R. Arnold, "Automated stereo perception," Dep. Comput. Sci., Stanford Univ., Stanford, CA, Tech. Rep. STAN-CS-83-961, Mar. 1983.
[9] H. Baker, "Depth from edge and intensity based stereo," Dep. Comput. Sci., Stanford Univ., Stanford, CA, Tech. Rep. STAN-CS-82-930, Sept. 1982.
[10] R. L. Henderson, W. J. Miller, and C. B. Grosch, "Automatic stereo recognition of man-made targets," *Soc. Photo-Opt. Instrum. Eng.*, vol. 186 (Digital Processing of Images), Aug. 1979.
[11] Y. Ohta and T. Kanade, "Stereo by intra and inter-scanline search using dynamic programming," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-5, 1983.
[12] R. Nevatia and K. R. Babu, "Linear feature extraction and description," *Comput. Graphics Image Processing*, vol. 13, pp. 257-269, 1980.
[13] R. S. Wallace, "A modified Hough transform for lines," in *Proc. IEEE Comput. Soc. Conf. Computer Vision and Pattern Recognition*, San Francisco, CA, June 1985.
[14] Y. Le Guilloux, "Determination automatique du mouvement dans une sequence d'images. Interet pour l'interpretation," Ph.D. dissertation, Ecole Nationale Supérieurre des Télécommunications, June 1984.
[15] M. J. Hannah, "Detection of errors in match disparities," in *Proc. Image Understanding Workshop*, Sept. 1982.

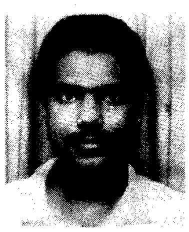

Rakesh Mohan (S'85–M'87) received the B.Tech. degree in computer science from the Indian Institute of Technology, Kanpur, India, in 1983.

He is currently a Research Assistant at the University of Southern California, Los Angeles, where he is working toward the Ph.D. degree in computer science. His research interests include computer vision, artificial intelligence, and artificial neural systems.

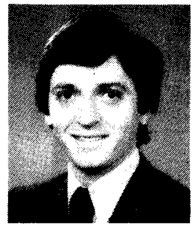

Gerard Medioni (S'82–M'83) received the Diplome d'Ingenieur Civil from the Ecole Nationale Superieure des Telecommunications, Paris, France, in 1977, and the M.S. and Ph.D. degrees in computer science from the University of Southern California, Los Angeles, in 1980 and 1983, respectively.

He is currently an Assistant Professor in the Departments of Electrical Engineering and Computer Science, University of Southern California. His research interests include computer vision, artificial intelligence, and robotics.

Dr. Medioni is a member of the Association for Computing Machinery and the American Association for the Advancement of Science.

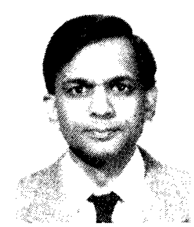

Ramakant Nevatia (S'71–M'74–SM'86) received the B.S. degree from the University of Bombay, India, and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, all in electrical engineering.

He has been with the University of Southern California, Los Angeles, since 1975, where he is currently Professor of Electrical Engineering and Computer Science and director of the Institute for Robotics and Intelligent Systems. He spent the academic year 1981–1982 at Stanford University as a Visiting Professor. He has authored two books: *Machine Perception* and *Computer Analysis of 3-D Curved Objects*. His research interests include computer vision, artificial intelligence, and robotics.

Dr. Nevatia is a member of the Association for Computing Machinery and the American Association for Artificial Intelligence. He is Associate Editor of the journals *Pattern Recognition* and *Computer Vision, Graphics, and Image Processing*, and the technical editor in the areas of robot vision and inspection systems for the IEEE JOURNAL OF ROBOTICS AND AUTOMATION.