

REPORT DOCUMENT		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER		2. REPORT CATALOG NUMBER	
AIM 1184		AD-A218 170	
4. TITLE (and Subtitle)		F. REPORT & PERIOD COVERED	
Stereo Feature Matching in Disparity Space		memorandum	
7. AUTHOR(s)		6. PERFORMING ORG. REPORT NUMBER	
David J. Braunegg			
9. PERFORMING ORGANIZATION NAME AND ADDRESS		8. CONTRACT OR GRANT NUMBER(s)	
Artificial Intelligence Laboratory 545 Technology Square Cambridge, MA 02139		N00014-85-K-0124 DACA76-85-C-0010	
11. CONTROLLING OFFICE NAME AND ADDRESS		10. PROGRAM ELEMENT PROJECT, TASK AREA & WORK UNIT NUMBERS	
Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209			
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE	
Office of Naval Research Information Systems Arlington, VA 22217		September 1989	
		13. NUMBER OF PAGES	
		20	
		15. SECURITY CLASS. (of this report)	
		UNCLASSIFIED	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report)			
Distribution is unlimited			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
DTIC ELECTE FEB 22 1990 S D D			
18. SUPPLEMENTARY NOTES			
None			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)			
Stereo Vision; artificial intelligence; algorithms; Stereo Matching; Disparity Space; (< >) ←			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)			
Abstract: This paper describes a new method for matching, validating, and disambiguating features for stereo vision. It is based on the Marr-Poggio-Grimson stereo matching algorithm which uses zero-crossing contours in difference of Gaussian filtered images as features. The matched contours are represented in <i>disparity space</i>, which makes the information needed for matched contour validation and disambiguating.			

(continued on back) →

Block 20 continued:

biguation easily accessible. The use of disparity space also makes the algorithm conceptually cleaner than previous implementations of the Marr-Poggio-Grimson algorithm and yields a more efficient matching process.

Keywords: continued previous page

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A. I. Memo No. 1184

September 1989

Stereo Feature Matching in Disparity Space

David J. Braunegg

Abstract: This paper describes a new method for matching, validating, and disambiguating features for stereo vision. It is based on the Marr-Poggio-Grimson stereo matching algorithm which uses zero-crossing contours in difference-of-Gaussian filtered images as features. The matched contours are represented in *disparity space*, which makes the information needed for matched contour validation and disambiguation easily accessible. The use of disparity space also makes the algorithm conceptually cleaner than previous implementations of the Marr-Poggio-Grimson algorithm and yields a more efficient matching process.

Acknowledgments. This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-85-K-0124 and under Army contract number DACA76-85-C-0010.

© Massachusetts Institute of Technology 1989

90 02 21 0541

1


The Marr-Poggio-Grimson stereo algorithm has been presented as a model of the human stereo vision process [Marr and Poggio] [Grimson 1981]. The algorithm was subsequently revised by [Grimson 1985] to improve its performance with no claims as to the relevance of these revisions to human stereo vision. Proceeding from this work, we present herein further modifications to the algorithm used to implement this theory of stereo vision. These modifications involve a change of representation which makes the constraints of the stereo matching problem more accessible and facilitates experimentation with methods of feature matching.

In the Marr-Poggio-Grimson stereo algorithm, the features which are matched between left and right images are zero-crossings of the images convolved with difference-of-Gaussian filters. The matching is performed across a range of eye fixation positions, with matches in a small area around each fixation position compressed into a single match which is then further processed. The method of feature matching presented here, rather than processing matches about individual fixation points, finds unique matches at each fixation point and then considers the set of matches across the full range of fixation positions.

2

The Marr-Poggio-Grimson stereo algorithm as implemented by Grimson consists of eight steps. These steps are outlined here; those unfamiliar with the algorithm are referred to [Grimson 1985] for details.

0. *Loop Over Levels:* Iterate from coarser to finer levels of representation. The level of representation is reflected by the value of the filter width w for the convolution step.
 1. *Convolutions:* Convolve the left and right images with $\nabla^2 G(w)$ filters.
 2. *Zero-Crossings:* Locate the non-trivial zero-crossings in the convolved images and mark these zero-crossings with their contrast signs.
 3. *Loop Over Fixation Position:* Iterate over a range of alignments of the left and right images.
 4. *Matching:* Identify valid feature matches between the left and right images.



- a. *Feature-Point Matching*: Given a disparity δ_0 , match positive zero-crossings against positive ones and negative zero-crossings against negative ones over a vertical range of $\pm\epsilon$ and a horizontal range of $\pm w$ about the current alignment.
 - b. *Figural Continuity*: Compress the contours matched about δ_0 from 4.a. into a single representation and eliminate those matched contours whose vertical lengths are less than a threshold.
 - c. *Disparity Updating*: For the remaining contours matched about δ_0 record in the disparity map for the current level the average disparity of the matched contour points.
5. *Loop*: Loop to Step 3. and repeat for all possible image alignments.
6. *Disambiguation*: Use zero-crossing matches from coarser channels to disambiguate matches at the current level.
7. *Loop*: Loop to Step 0. and repeat for next finer level of representation.
8. *Consistency*: Eliminate zero-crossing matches which are inconsistent with coarser channel matches.

We propose a modification to the above algorithm at Steps 4 and 5. In the modified version, feature points are not matched in a range about each fixation position, but rather are matched only at the fixation positions. Figural continuity is then applied to the matched contours as they lie across the full range of disparities (fixations) considered. (A figural continuity test was later added to Grimson's implementation as well [Grimson 1989].) Specifically, the new steps of the algorithm are

- 4. *Matching*: Identify valid feature matches between the left and right images.
 - a. *Feature-Point Matching*: Given a disparity δ_0 , match positive zero-crossings against positive ones and negative zero-crossings against negative ones over a vertical range of $\pm\epsilon$ about the current alignment. The $\pm\epsilon$ is used solely to allow for a (possible) slight vertical misalignment between the images.
 - b. *Disparity Updating*: For the contours matched at δ_0 record in the disparity map for the current level the disparity of the matched contour points.
- 5. *Loop*: Loop to Step 3. and repeat for all possible image alignments.

- 5-1. *Contour Following*: Form linked lists of matched contour points in the disparity map for the current level. These linked lists are the candidate matched contours for their corresponding left-image contours.
- 5-2. *Contour Validity Checking*: Test the matched contours in the current level's disparity map for validity.
 - a. *Horizontal Segment Extension*: Connect matched contour segments across regions of horizontal contour points where the horizontal extensions have a disparity gradient less than a threshold, h .
 - b. *Figural Continuity*: Eliminate those matched contours whose vertical lengths are less than a threshold, f . *The contours are followed without regard to how they vary in disparity.*
 - c. *Disparity Gradient*: Eliminate segments of matched contours which vary in disparity greater than a threshold, g . Recheck figural continuity for any contour from which a segment is removed.
- 5-3. *Matched Contour Disambiguation*: Disambiguate contour points with two or more candidate matches.
 - a. *Contour Subsumption*: Eliminate those contour matches which are subsumed by other matches.
 - b. *Consistent Disparities*: Where two or more matches remain for a contour point, choose the disparity most consistent with the disparities found for the rest of the contour.

The changes made to the algorithm will be discussed in the following sections. At this point we simply note that the differences lie in how the contours are matched and how the matched contours are checked for validity. The validity checks listed above are modifications of tests from the original algorithm while the disambiguation tests are new.

3 Disparity Space Contours

The new steps in the stereo matching algorithm involve a change in representation. In the new algorithm, matched contours are represented as contours in *disparity space*. *Disparity space* is a three-dimensional space: the x - and y -dimensions are the same as in one of the images (typically the left image) (Figure 1) while the third dimension is disparity. Matched contours are plotted as sets of ordered points in this space. Thus, for any matched pair of contour

points (Figure 2), the matched point is explicitly represented in terms of its position in the (left) image and its perceived disparity (Figure 3). If we were to convert the disparity measurements to distance from the cameras (using information about the camera geometry), the disparity space representation would become a scaled model of the arrangement of the physical contours in the three-dimensional world.




Figure 1. Left and right images.




Figure 2. Left- and right-image contours. A single left-image zero-crossing contour is shown. For the right image the area of possible matches (in dotted lines) for the left-image contour as disparity varies is shown. Also shown are the right-image contours which are candidate matches for the chosen left-image contour.

When performing operations on the matches of individual left-image contours, we do not need to deal with them in the full three-dimensional disparity space.




Figure 3. Disparity Space, showing the candidate matches for the left-image contour that appears in the dotted box of Figure 2.

Since we know where a contour lies in the x and y dimensions (the left image location of the contour), we can consider the matches of this contour in a two-dimensional plane—the first dimension being distance along the contour and the second disparity (Figure 4). This *disparity-space plane* is a two-dimensional plane embedded in the three-dimensional disparity space.

It follows that points of the contour matches are referenced by arc length (based at the beginning of the contour) and disparity rather than x -coordinate, y -coordinate, and disparity. Thus, three-dimensional operations on contour matches in disparity space are reduced to two-dimensional operations in the disparity-space planes.




Figure 4. Disparity-Space Plane for the left-image contour of Figure 2, showing the candidate matches. This is the two-dimensional plane embedded in the three-dimensional disparity space shown in Figure 3.

The disparity space representation for matched contour points has several benefits. By considering a matched contour as a whole, this new representation aids in the selection of breakpoints for long contours (Section 4.1) and the extension of matches across horizontal contour segments (Section 5.1). (Long contours often exist in a difference-of-Gaussian filtered image due to the requirement that zero-crossings form closed contours rather than due to the existence of corresponding real-world object features.) The representation allows us to implement a cleaner form of the figural continuity constraint (Section 5.2). It facilitates the application of a disparity gradient constraint to the whole matched contour or any part thereof (Section 5.3). Finally, the disparity space representation allows us to easily check for consistent disparities and contour subsumption when disambiguating candidate contour matches (Sections 6.1 and 6.2).

4 Matching to Yield Disparity Space Contours

Following from [Grimson 1985], we assume that the convolved left and right images,

$$LC_w(x, y) = \nabla^2 G(w) * L$$

$$RC_w(x, y) = \nabla^2 G(w) * R$$

have been computed. $\nabla^2 G(w)$ denotes the Laplacian of a Gaussian whose central negative portion has width w , $*$ is the convolution operator, and L and R denote the left and right images, respectively. For each of these convolved images, the nontrivial zero-crossings have been located and marked with their contrast signs, yielding the bit maps:

$$LP_w(x, y) = \text{positive zero-crossings of } LC_w(x, y)$$

$$LN_w(x, y) = \text{negative zero-crossings of } LC_w(x, y)$$

$$LH_w(x, y) = \text{horizontal zero-crossings of } LC_w(x, y)$$

$$RP_w(x, y) = \text{positive zero-crossings of } RC_w(x, y)$$

$$RN_w(x, y) = \text{negative zero-crossings of } RC_w(x, y)$$

$$RH_w(x, y) = \text{horizontal zero-crossings of } RC_w(x, y).$$

Now, to find the candidate right-image matches for the k^{th} left-image contour $LC_{w_0, k}$ at the current filter width w_0 , we construct a disparity-space plane

$$P_{w_0, k}(\delta, i) = \begin{cases} RP_{w_0}(X, Y) & \text{for } LC_{w_0, k}(i) \in LP_{w_0} \\ RN_{w_0}(X, Y) & \text{for } LC_{w_0, k}(i) \in LN_{w_0} \\ RH_{w_0}(X, Y) & \text{for } LC_{w_0, k}(i) \in LH_{w_0} \end{cases}$$

over

$$\delta_{\min} \leq \delta \leq \delta_{\max} \quad (\text{disparity range}),$$

$$0 \leq i < \text{length}(LC_{w_0,k}),$$

where

$X = X(E(w_0, k, i, \delta))$ is the x-coordinate of $E(w_0, k, i, \delta)$,

$Y = Y(E(w_0, k, i, \delta))$ is the y-coordinate of $E(w_0, k, i, \delta)$,

$E(w_0, k, i, \delta)$ is the point in the left image plane δ pixels to the right
along the epipolar line through $LC_{w_0,k}(i)$,

$LC_{w_0,k}(i)$ is the i^{th} point in contour $LC_{w_0,k}$.

In the case of horizontal epipolar lines, matching the points of horizontal contours is ambiguous. For such points, we rely on the Horizontal Segment Extension of Step 5-2.a. for determination of the disparities. The disparity-space planes then become

$$P_{w_0,k}(\delta, i) = \begin{cases} RP_{w_0}(X, Y) & \text{for } LC_{w_0,k}(i) \in LP_{w_0} \\ RN_{w_0}(X, Y) & \text{for } LC_{w_0,k}(i) \in LN_{w_0} \\ 0 & \text{for } LC_{w_0,k}(i) \in LH_{w_0} \end{cases}$$

Certain computational advantages are realized by implementing the disparity space contour representation for matching zero-crossings instead of the $\pm w$ window method of [Grimson 1985]. At a given disparity δ_0 , only one point in the right image must be compared with a contour point in the left image. Grimson's method entails combining $2w + 1$ points from the right image into a single data point and comparing this data point with a contour point in the left image. If we assume horizontal epipolar lines in the images, the fact that we only match points at a single horizontal disparity allows us to search for matching points by simply scanning across the right image over the range of horizontal disparities (see Figure 2). No additional processing is required to consider $\pm w$ windows around each fixation position. (We note that in both algorithms the $\pm \epsilon$ vertical range is achieved by overlaying copies of the right image at the required vertical offsets and using the resulting image for matching.)

4.1 Contour Following in Disparity-Space Planes

Once we have formed the disparity-space plane for a left-image contour, the candidate matching contours from the right image are easily found (Step 5-1.). We simply scan the rows of the disparity-space plane until we find a matched point. We then trace matched points along the contour starting from this matched point, forming a linked list of matched points. Given matched point

$P_{w_0,k}(\delta_0, i_0)$, we search for the next matched point in the range $[P_{w_0,k}(\lfloor \delta_0 - c \rfloor, i_0 + 1), P_{w_0,k}(\lceil \delta_0 + c \rceil, i_0 + 1)]$ where c is a threshold on the maximum jump in disparity between successive contour points (larger than that allowed by the disparity gradient discussed in Section 5.3). All of the candidate matching contours in the disparity-space plane are found in this way, possibly breaking the original contour into several matched contour pieces. We then check these candidate matches for validity and disambiguate when multiple matches have been found for the left-image contour.

As stated in Step 4.a., we consider candidate matches over a vertical range of $\pm \epsilon$ about the epipolar line along which we expect to find matches. Because of this, the matched points in a disparity-space plane may form lines that are more than one pixel wide. We perform a thinning operation on such lines to obtain one pixel wide candidate matching contours in the disparity-space plane. Rather than using a standard "grass fire" approach to thinning, we bias the thinning algorithm to prefer lines which vary least in disparity over their lengths. Due to this thinning operation and the fact that the y-axis of the disparity-space plane represents arc length along the contour, the resulting candidate matches in the plane form lines which are strictly monotonic in the vertical axis of the plane.

5 Matched Contour Validity Checking in Disparity Space

After obtaining the candidate matching right-image contours for each left-image contour, we apply several validity tests to each of the candidate contours in disparity space. These tests are applied to the candidate contours individually and must be applied to every candidate contour. We can not assume that a candidate match is valid simply because it is a unique match for a particular left-image contour.

5.1 Horizontal Segment Extension

Contour matching can not be applied to horizontal contour segments when we have horizontal epipolar lines. However, if we do not consider horizontal points, we may unintentionally break a long matched contour segment into two or more shorter contour segments. To avoid this problem, we extend contour segments across horizontal sections where possible (Step 5-2.a.). Suppose that a matched contour segment in a disparity-space plane ends at the point

$P_{w_0,k}(\delta_0, i_0)$. We search the points in the range $[P_{w_0,k}(\lfloor \delta_0 - jh \rfloor, i_0 + j), P_{w_0,k}(\lfloor \delta_0 + jh \rfloor, i_0 + j)]$ for $j = 1, 2, 3, \dots, \text{length}(LC_{w_0,k})$ until we find a matched contour point. (h is the threshold mentioned in Step 5-2.a. of the modified algorithm.) If no point is found, we cannot extend the contour. If a point is found, say at $j = j_n$, then the points

$$P_{w_0,k}(\delta_0 + 1, i_0 + 1), P_{w_0,k}(\delta_0 + 2, i_0 + 2), \dots, P_{w_0,k}(\delta_0 + j_n, i_0 + j_n)$$

are added to the contour and we continue to trace the contour in the disparity-space plane from the point $P_{w_0,k}(\delta_0 + j_n, i_0 + j_n)$.

Typically, the threshold h is set to the same value as the threshold c mentioned in Section 4.1 above. The threshold h is the maximum allowable jump in disparity between contour points where one or both of the points comes from a horizontal section of the contour.

Horizontal segment extension is facilitated by the disparity space representation of candidate matches (Figure 5). To extend a match across a horizontal section of a contour, we simply search for the next y-coordinate (in the direction of increasing arc length) until matched points are found. Choose the matched point at this y-coordinate whose x-coordinate (disparity) is closest to that of the original point. If any y-coordinates were crossed which correspond to non-horizontal contour points, the extension fails. If we have extended the contour for n points and the difference in disparity across the extension does not exceed nh , then we accept the extension and connect the matches across that horizontal contour section.




Figure 5. Contour in a disparity-space plane before and after horizontal segment extension.

5.2 Figural Continuity

Arguments based on both the cohesiveness of matter and psychophysical evidence support the use of figural continuity as a verification criterion for stereo matches [Mayhew and Frisby, 1980] [Mayhew and Frisby, 1981]. In the implementation of the figural continuity constraint given by [Grimson 1985], one pixel wide gaps were allowed in a contour when checking its length. After experimenting with the stereo algorithm, however, we have found that this allowable gap feature is rarely used. Therefore, we have implemented a simplified figural continuity constraint based on the continuous length of a matched contour (Step 5-2.b.). To apply the figural continuity test, we simply eliminate all matched contours which are vertically shorter than f pixels long.¹ Contour points from horizontal segments are not counted in the length of a contour because they are interpolated between matched contour points instead of resulting from matched points themselves.

Note that the length of the contour is determined regardless of how the contour varies in disparity. The Disparity Gradient test (see the following section) explicitly considers how a contour varies in disparity. The disparity space representation of a matched contour allows the figural continuity constraint to be applied to the contour as a whole. In the original algorithm, the length of a matched contour can only be checked in a window of width $2w$ pixels of disparity. As presented by Mayhew and Frisby, figural continuity should be applied to a matched contour as a whole, not just to arbitrary pieces of it. With the disparity space representation for matched contours, we have separated the length requirement for matched contours from consideration of how widely they range in disparity.

5.3 Disparity Gradient

Psychophysical observations have shown the importance of a disparity gradient limit for stereopsis [Burt and Julesz, 1980a], [Burt and Julesz, 1980b]. These observations were the basis of the PMF stereo algorithm [Pollard, Porrill, Mayhew, Frisby], [Pollard, Mayhew, and Frisby]. However, only a weak condition on the disparity gradient is incorporated in the original Marr-Poggio-Grimson stereo algorithm. Below we provide an explicit check of the disparity gradient using the disparity-space representation of contour matches.

Disparity gradient for a matched contour is directly interpreted from the slope of the matched contour in the disparity-space plane for the contour. (The x-axis corresponds to disparity and the y-axis to arc length along the

¹The threshold f is the *figural continuity length* and is typically expressed as a fraction of the width of the central region of the $\nabla^2 G$ filter.

contour.) Since the discrete nature of the representation makes defining and finding gradients difficult, we define a procedure for finding the gradient of a matched contour over various sections of the contour (Step 5-2.c.). Rather than choose a fixed segment length over which to calculate the gradient, we let the segment length vary.

To apply the Disparity Gradient constraint to a matched contour, we first obtain a straight-line approximation [Pavlidis] to the contour in the disparity-space plane, including the horizontal segment extensions, if any (Figure 6). We let points deviate from these straight-line segments by two pixels (typically) to allow for positional errors due to discretization and the $\nabla^2 G$ convolution. Segments with disparity gradient greater than g are deleted, while the remaining segments are accepted as valid contour matches by this processing step.




Figure 6. Straight-line approximation to a matched contour in a disparity-space plane.

It is important to note that Step 5-2.c. (Disparity Gradient) of the new algorithm was implicit in Step 4.a. of the original algorithm. In the original algorithm, a matched contour was not permitted to vary by more than $2w$ pixels across the figural continuity length. It could, however, vary quite rapidly within these bounds (and thus contain large disparity gradients) and still be accepted. The new algorithm places a specific limit on how quickly a contour may vary in disparity (i.e., a disparity gradient limit). This disparity gradient test is distinct from the figural continuity test and also does not affect the matching process since we match contours at a particular horizontal disparity instead of over a range of disparities.

The fact that we explicitly check the disparity gradient along a matched contour allows us to break the contour into distinct segments, eliminating

parts of the contour whose disparity gradients are too steep. This situation often occurs on zero-crossing contours in $\nabla^2 G$ images since these contours must be closed. Many times strong zero-crossing contour segments are closed by weak zero-crossings which do not correspond to physical features in the real world. Due to the nature of these weak zero-crossing segments, they tend to wander in the image, producing matched zero-crossing segments which vary in disparity. The varying disparity of these weak contour segments is typically evidenced by high disparity gradients, allowing many of these segments to be located and eliminated.

Typically, the threshold g is smaller than the thresholds c and h used in Steps 5-1. and 5-2.a. (see Sections 4.1 and 5.1 above). This is due to the fact that a contour can vary sharply in disparity over a few points but still have a low disparity gradient overall.

6 Matched Contour Disambiguation

After validating the matched contours using the figural continuity and disparity gradient criteria, ambiguous matches often remain. We attempt to resolve these ambiguities by checking for contour subsumption and consistent disparities along contours. Afterward, any remaining ambiguities are resolved, if possible, by considering nearby contour matches from the current and coarser filter channels (as in the original Marr-Poggio-Grimson stereo algorithm).

6.1 Contour Subsumption

When two or more matches are possible for a given contour, one of the matches frequently subsumes the others. That is, one candidate match extends over a larger portion of the contour and the portion that it does cover includes some of the other matches. If this is the case, we accept the subsuming match as the correct one (Step 5-3.a.). This is justified because the base of support for that match is stronger (hence the longer length of the match) than for the others.

The disparity space representation of contour matches facilitates direct comparison of the extents of matches of consistent disparities (Figure 7). The starting and ending points of the matched portions of contours are directly represented in disparity space. In our implementation of contour subsumption disambiguation, we initially require strict subsumption, i.e., the subsuming contour must completely cover the range of the subsumed contour as well

as extend above and below it. This requirement is relaxed if the subsuming contour extends well beyond one end of the subsumed contour but falls a few pixels short of the other end.

Thus, the candidate match

$$P_{w,k}(\delta_i, i), P_{w,k}(\delta_{i+1}, i+1), \dots, P_{w,k}(\delta_{i+m}, i+m)$$

for left-image contour k subsumes candidate match

$$P'_{w,k}(\delta_j, j), P'_{w,k}(\delta_{j+1}, j+1), \dots, P'_{w,k}(\delta_{j+n}, j+n)$$

if either

$$i < j \text{ and } i + m > j + n,$$

$$i < j - [(j + n) - (i + m)] \text{ and } i + m < j + n,$$

or

$$i + m > j + n + [(i + m) - (j + n)] \text{ and } i + m > j + n.$$




Figure 7. Contours in a disparity-space plane before and after subsumed contours are removed.

6.2 Consistent Disparities

To disambiguate remaining contour points with more than one match, we choose the match whose disparity is most consistent with the disparities found for the rest of the contour (Step 5-3.b.) (Figure 8). Consider a contour point with matches at disparities $\delta_1, \delta_2, \dots, \delta_n$. Let N_{δ_i} be the number of points from the same contour that are unambiguously matched at disparity δ_i . Assign disparity δ_j to the contour point when δ_j maximizes N_{δ_j} for $1 \leq j \leq n$ and $N_{\delta_j} > 0$. If all of the $N_{\delta_j} = 0$, i.e., there are no unambiguous contour points whose disparity matches one of the possible disparities

of the point in question, then select the match whose disparity is closest to an unambiguous disparity for the contour: choose δ_j such that δ_j minimizes $|\delta_j - \delta_i|$ for $1 \leq j \leq n$ over all i . However, if this minimum distance exceeds a threshold, i.e., $\min |\delta_j - \delta_i| > \epsilon, \forall i, j$, then the ambiguously matched point is left unmatched at this stage. These remaining ambiguous points will be disambiguated by using matched contour information from coarser channels (Step 6).




Figure 8. Contours in a disparity-space plane before and after ambiguous matches with inconsistent disparities are eliminated.

7 Advantages of the Disparity-Space Approach

The use of the disparity-space plane representation for candidate contour matches has several advantages over the original Marr-Poggio-Grimson implementation. These advantages are based on the direction representation of arc length vs. disparity for the candidate matches of a contour.

With the disparity-space plane representation for the contour matches, the figural continuity is directly observable from the candidate matches. Using this representation, the figural continuity constraint is applied once per matched contour. The original implementation required this constraint to be checked for the matches found in the $\pm w$ range at each fixation position. Even assuming that a matched contour lies at a single disparity, the original method would check the figural continuity of this contour $2w + 1$ times.

To check the disparity gradient of a matched contour, we must follow the contour in the x- and y-coordinates of the original image as well as in

disparity. However, by representing contour matches in disparity-space planes, we consider the matched contours in a two-dimensional space—one dimension is disparity while the second is distance along the contour. Thus contour points are referenced by disparity and arc length (based at the beginning of the contour) rather than disparity, x-, and y-coordinates (Figure 4). We can therefore compare candidates for zero-crossing matches in these disparity-space planes instead of in a three-dimensional space.

Contour subsumption is easily determined in the disparity-space planes. The linked lists of points which comprise a match for a contour are parameterized by arc length, allowing direct comparison of different candidate matches.

Finally, consistency of disparity along a contour can be enforced naturally using the disparity-space plane representation. This follows from the fact that the matched contour points are parameterized by disparity in the representation.

8 Conclusions

By considering matched contours in disparity space, we are able to improve the Marr-Poggio-Grimson in several ways. The figural continuity constraint can be applied to complete matched contours instead of sections of them which are bounded by a fixed disparity range. An explicit disparity gradient threshold can be applied to matched contours and sections of the contours which do not meet this constraint can be removed. Computational savings are realized by eliminating the horizontal disparity window formerly used in the matching process (Step 4.a.), both by allowing a simpler matching process and by permitting us to apply the figural continuity constraint only once per matched contour.

Disambiguation of competing contour matches is facilitated by the disparity-space plane representation of the matches. The explicit representation of arc length makes contour subsumption easy to determine. The explicit representation of disparity along the contour matches allows us to check the consistency of the disparity of candidate matches in a straightforward manner.

With this new disparity-space plane representation for stereo matching, the task of the matching problem is obvious—the best matches appear as long, connected lines in the disparity-space plane. The explicit representation of arc length and disparity of the candidate matches aids validation and disambiguation of the matches. Finally, the unnecessary detail of the actual location of

the contours is eliminated for the matching, validation, and disambiguation operations.

9 References

- Burt, P. and B. Julesz, "A Disparity Gradient Limit for Binocular Fusion," *Science*, 208, 1980a, 615-617.
- Burt, P. and B. Julesz, "Modifications of the Classical notion of Panum's Fusional Area," *Perception*, 9 (6), 1980b, 671-682.
- Grimson, W. E. L., "Computational Experiments with a Feature Based Stereo Algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-7 (1), January 1985, 17-34.
- Grimson, W. E. L., *From Images to Surfaces: A Computational Study of the Human Early Visual System*, MIT Press, Cambridge, MA, 1981.
- Grimson, W. E. L., personal communication, 1989.
- Marr, D. and T. Poggio, "A Theory of Human Stereo Vision," *Proceedings of the Royal Society of London*, B (204), 1979, 301-328.
- Mayhew, J. E. W. and J. P. Frisby, "The Computation of Binocular Edges," *Perception*, 9, 1980, 69-86.
- Mayhew, J. E. W. and J. P. Frisby, "Psychophysical and Computational Studies towards a Theory of Human Stereopsis," *Artificial Intelligence*, 17, 1981, 349-385.
- Pavlidis, T., *Algorithms for Graphics and Image Processing*, Computer Science Press, Rockville, MD, 1982, 287-290.
- Pollard, S. B., J. Porrill, J. E. W. Mayhew, and J. P. Frisby, "Disparity Gradient, Lipschitz Continuity, and Computing Binocular Correspondences," *Third International Symposium on Robotics Research*, Gouvieux, France, O. D. Faugeras and G. Giralt (eds.), MIT Press, Cambridge, MA, 1985, 19-26.
- Pollard, S. B., J. E. W. Mayhew, and J. P. Frisby, "PMF: A Stereo Correspondence Algorithm Using a Disparity Gradient Limit," *Perception*, 14, 1985, 449-470.

Appendix

Although the simplified figures shown in the body of this paper help one to understand the disparity-space representation and operations, it is interesting to see the effects of these operations on a real image. Unfortunately, the limits of black-and-white reproduction make it difficult to present the depth results of the stereo algorithm. Since this paper addresses the problems associated with finding the correct matches for image contours, these contours are shown below through the various stages of matching, validation, and disambiguation for a real image.



Figure 9. The left and right images of a stereo pair.




Figure 10. Original unambiguously matched left contour points (Step 4.).




Figure 11. Unambiguous left contour matches after horizontal segment extension (Step 5-2.a.).




Figure 12. Left contour matches after imposing figural continuity constraint (Step 5-2.b.).




Figure 13. Left contour matches after imposing disparity gradient constraint (Step 5-2.c.).



Figure 14. Left contour matches after eliminating subsumed contours (Step 5-3.a.).



Figure 15. Left contour matches after eliminating inconsistent disparities (Step 5-3.b.).