[3] Implementation Details of the PMF Stereo Algorithm

Stephen B Pollard, John E W Mayhew and John P Frisby

AI Vision Research Unit,
University of Sheffield, Sheffield S10 2TN, UK

INTRODUCTION

This paper describes the current version of the PMF stereo algorithm. The two preceding papers (Pollard, Mayhew and Frisby, 1985 [1]; Pollard, Porrill, Mayhew and Frisby, 1985 [2]) were largely concerned with the theoretical considerations that have underpinned the general design philosophy of PMF. Within that general framework a great deal of scope exists for particular implementation details. The ones described here attend to the twin requirements of robustness and efficiency on state of the art computer machinery. Also, this implementation incorporates some features from other stereo algorithms aimed at exploiting more global constraints than those embedded in the original PMF.

CONTROL STRUCTURE

Figure 1 provides an overview of the control structure. Processing stages are represented as rectangular boxes, and decisions as diamonds. The algorithm makes three iterations through a fixed sequence of processing stages. Once selected, matches remain fixed for any remaining iterations. For selection, matches must satisfy a number of matching criteria based upon gradient support, figural continuity, and the satisfaction of the ordering constraint.

POTENTIAL MATCH SELECTION

The physical locations of potential matches satisfy the epipolar constraint. In particular, it is assumed that the edges to be matched have arisen from a simplified parallel camera geometry in which the imaging planes of each camera are coplanar, the inter-ocular axis parallel to it, and the images arrays similarly oriented. Under such conditions the epipolar constraint becomes a same-raster constraint, thereby simplifying the process of match selection. Whilst this restriction may seem severe, it is in practice straightforward to rectify edge locations to give the appearance that they arose from a parallel camera geometry provided that the true camera geometry is known to reasonable accuracy. Thus rectification assumes either a priori knowledge or some form of calibration process to determine the physical camera geometry. In the examples of the performance of PMF shown below we have used the method of calibration from Tsai (1986) to recover a suitable approximation to the physical camera geometry.

Initial ranges of allowable image disparity can be either (1) preset arbitrarily: for later performance examples they were fixed to +/- half the size of the images; or (2) selected using disparity histogramming (based upon Shirai and Nishimoto, 1986) to set initial disparity ranges for each region of the image (16 by 16).

Matches between edges are restricted to have differences in left and right image orientations that could have arisen from a disparity gradient (DG) of 1.0. Hence edges that have orientations close to vertical are allowed to reorient to a greater extent than those close to horizontal. The sign and physical contrast of a pair of edges can also be used to determine both their matchability and matching strength. In the initial stages of the algorithm the contrast of a pair of matching edges must agree to a factor of 3. The strength of a match is the product of the individual contrast values, thus giving greater weight to matches between more robust edges.

SELECT SEED POINTS

A strategy aimed at fast yet reliable matching is to begin by seeking a set of strong seed point matches which can be used to guide selection of subsequent matches. This is done by trying to match at first only a subset of all the edge points available. So, only every nth point on an edge string is considered (currently n=4) as a seed point, and only every mth point along a string can give support (currently m=2). Figure 2 illustrates these concepts.

COMPUTE WITHIN-DG-LIMIT SUPPORT

In figure 3 the small circles $p, p', p'', i,$ and j' represent edge point primitives (the edge strings of which they form a part are not shown). Matching is initiated from left image primitives and two potential matches for p are illustrated with the lines labelled pp' and pp''. Support for each of these matches is sought from the potential matches of other left image primitives lying within the neighbourhood support circle (of radius r not drawn to scale) shown around p. Just one supporting match ij' is illustrated for the match pp': note that ij' lies within the disparity gradient limit with respect to pp'.

The local neighbourhood support scheme is run with the DG limit set to 0.5, a value which provides better disambiguating power than 1.0 (Pollard, 1985). The size of the support neighbourhood is set to be a function of the image size (currently a default radius of 20 pixels for 256x256 images, a radius of 40 pixels for 512x512 images). The underlying concern here is to ensure that the support neighbourhood is large enough to provide good disambiguation while not being so large that processing time is spent needlessly.

ENFORCE WITHIN-DG-LIMIT SELECTION

From the list of potential matches for a given left image seed point candidate, choose the strongest match *unless* there exists in the neighbourhood a point with a stronger match that exceeds DG=1.5 with respect to the match under consideration. If a selection is made, eliminate from further consideration at this stage (they may get reconsidered later): (i) all other matches for the given left image edge primitive, and (ii) those that violate the gradient limit with respect to the selected match.

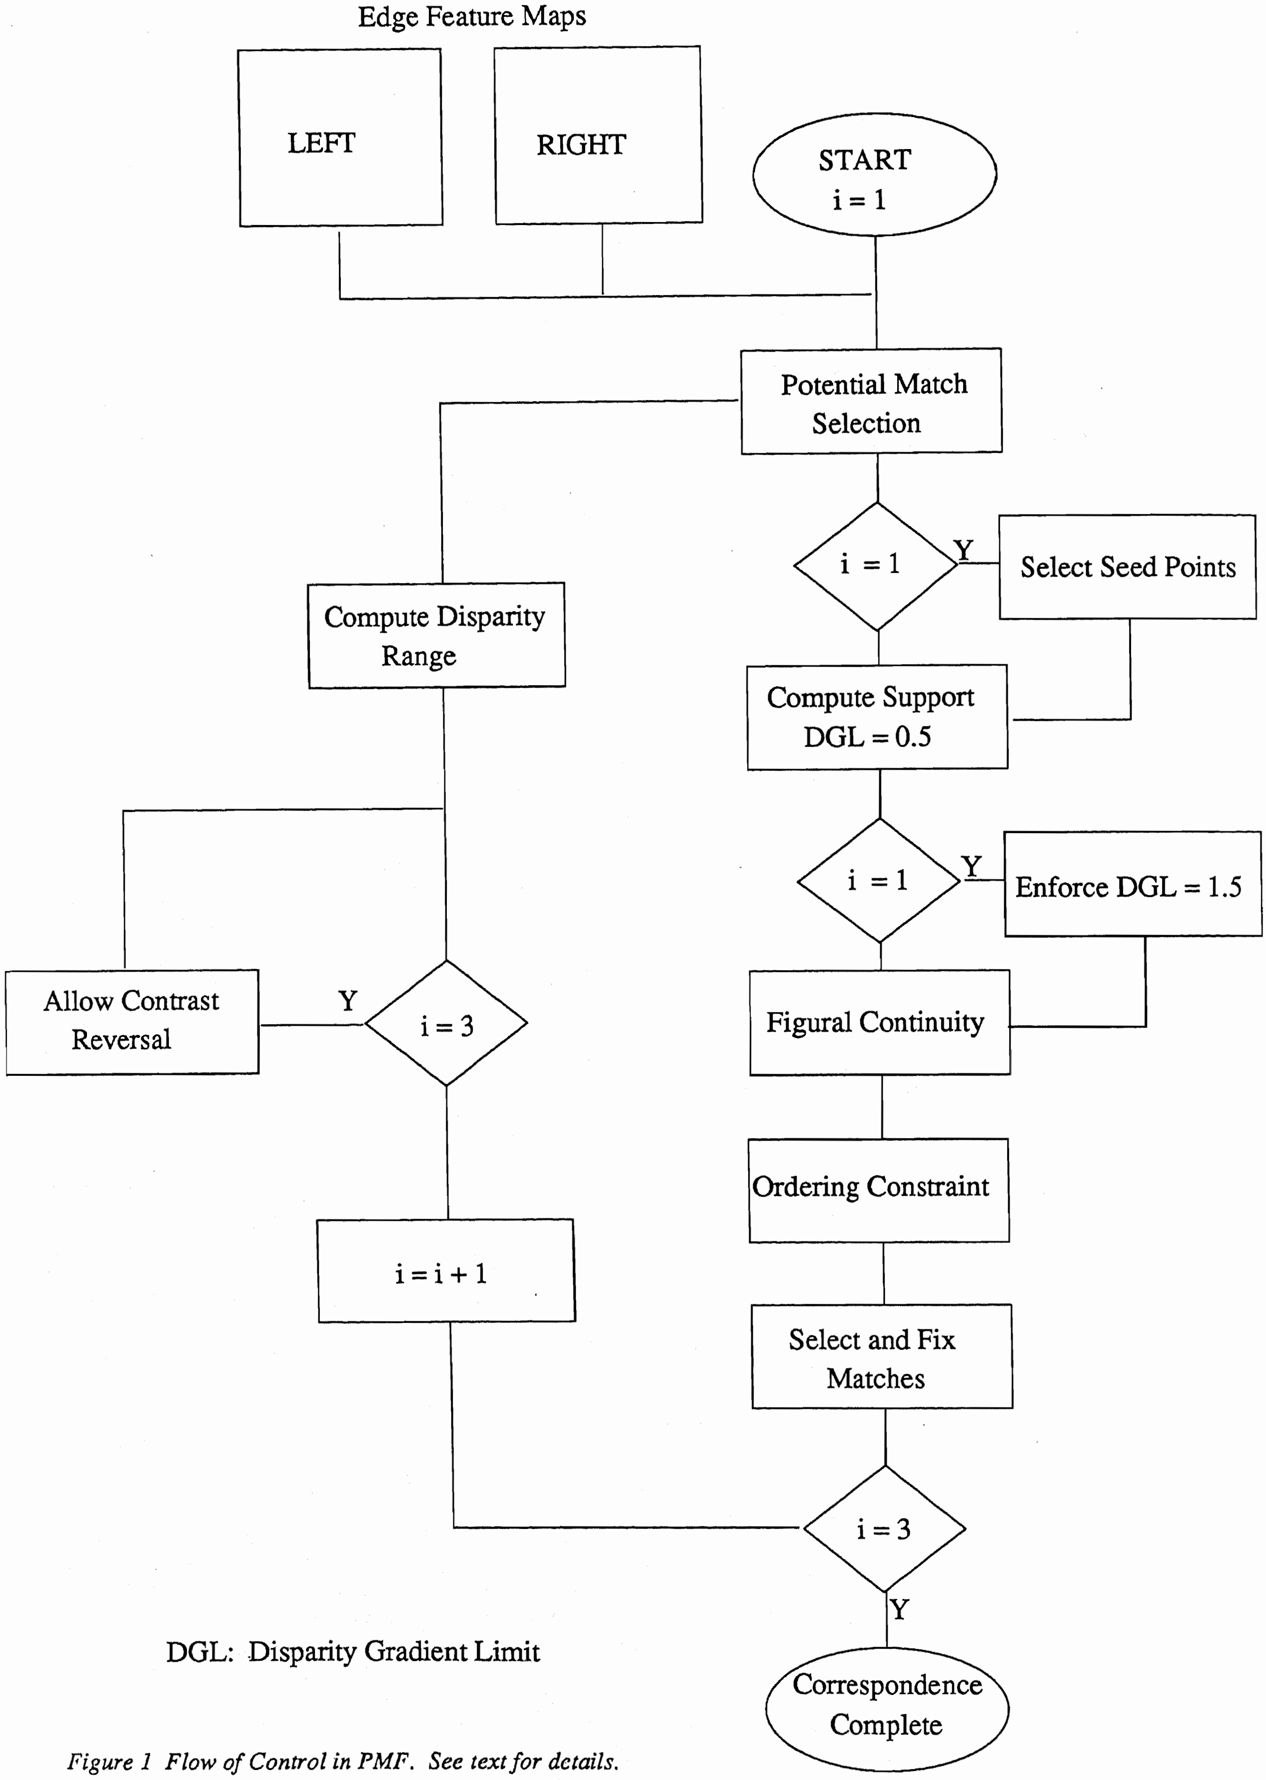

DGL: Disparity Gradient Limit

Figure 1 Flow of Control in PMF. See text for details.

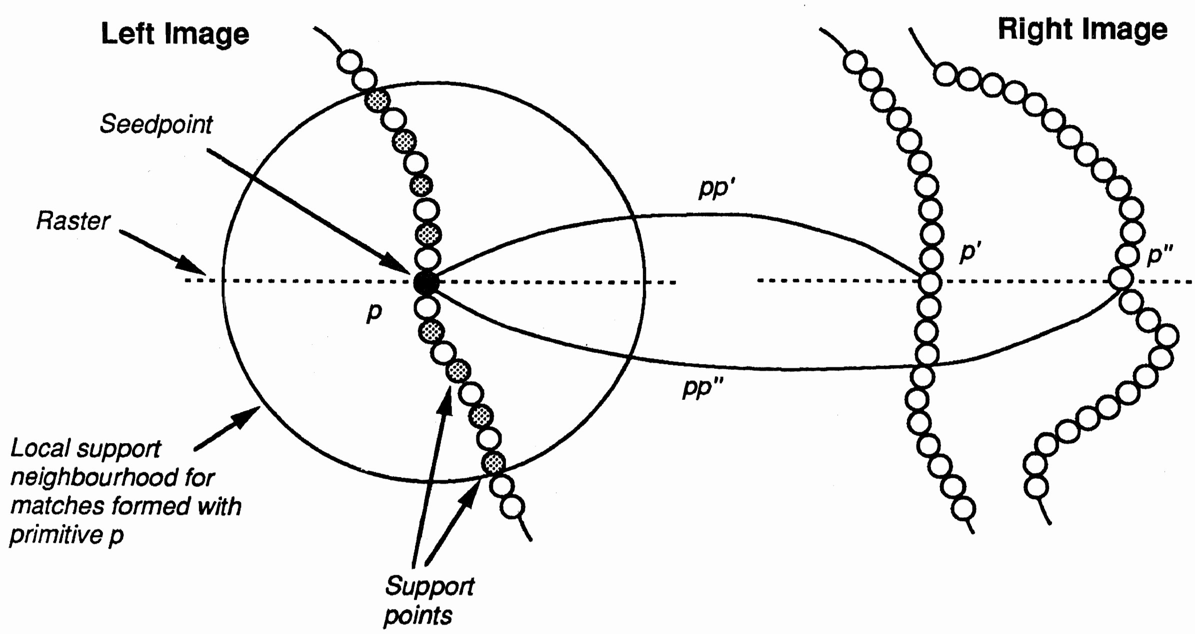

Figure 2 Selection of Seed Point Matches. The small circles represent edge point primitives comprising edge strings. Matching is initiated from left image strings and at first only every nth point comprising a left hand string is considered for matching. Such points are termed 'seed points'. Two potential matches for the seed point primitive p are shown, pp' and pp". Support for each of these matches is sought within the neighbourhood support circle shown around p (radius not drawn to scale), with every mth primitive along the string evaluated for the support it offers (the dotted circles illustrate m=2).

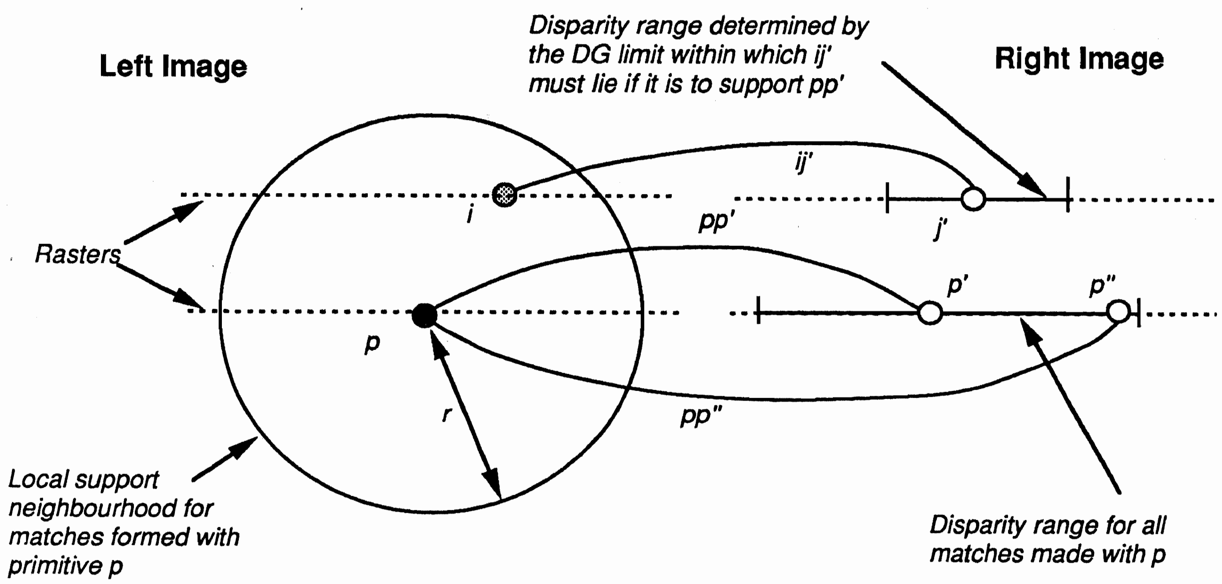

Figure 3 Neighbourhood Support. See text for details.

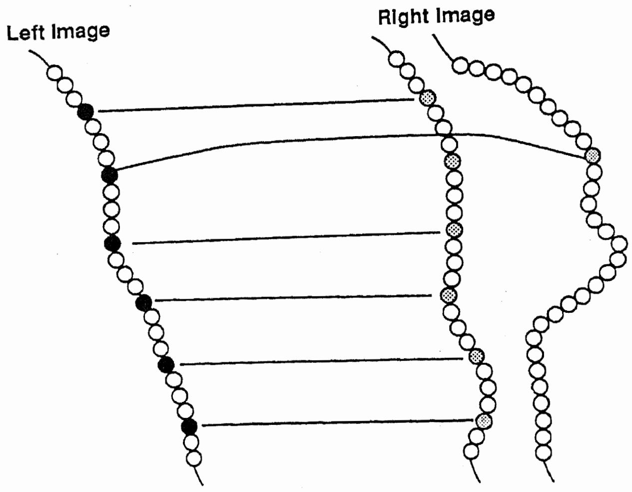

Figure 4 Exploiting the Figural Continuity Constraint by Matching Strings of Edge Points as a Whole. Lines connecting left and right image strings depict matches between initial seed points. The second line from the top shows a match whose right image primitive lies within a different edge string from those of all other matches. This fact is discovered and the discrepant match killed off.

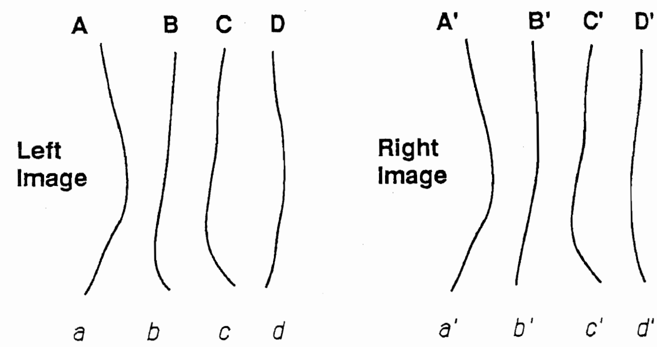

Figure 5 Imposing the Ordering Constraint Explicitly. Suppose strings of points existed in the two images as shown, each with associated strengths a, b, ... c', d'. Suppose also that the correct matches were AA', BB', CC' and DD'. If matches AA', BD', CC' and DB' were initially chosen then the violation of the ordering constraint by matches BD' and DB' would be noted. The weakest strings whose elimination from matching would remove the violation are then discovered, and their primitives freed for reconsideration by later stages of matching.

The rationale for the *unless* proviso is that if there exists a stronger match in the neighbourhood that offends a quite steep disparity gradient limit then this is evidence that the match under consideration might well be false. It should not therefore be selected as a seed point on which to base other selections.

The selection rule is applied iteratively. If a stronger match exists then no selection is made on that iteration. If the stronger match gets killed off (by not getting selected as a seed point itself) then any selection held over may be allowed to proceed on the next iteration.

The above selection procedure implements the uniqueness constraint by allowing only one match per primitive but it does so only with respect to the left image points. Uniqueness is not imposed with respect to right image points as well (as happened in early versions of PMF) except later on via a greater use of the ordering constraint

FIGURAL CONTINUITY

The figural continuity constraint is exploited by counting the seed points along each string of edge points, and selecting that string which contains the majority of seed point matches. This operation marks a shift from the fairly local operations embedded in the neighbourhood support stage to much more global operations that can in principle span quite large regions of the image (in fact the whole image if an edge happens to traverse right across it). Seed point matches to the 'wrong' edge string get killed off at this stage (figure 4).

This stage copes with mild 'wallpaper illusion' problems caused by similar but not identical edge strings. Of course, there is no way to solve the ambiguity problem posed by repeating identical surface texture elements (the *real* wallpaper illusion problem) unless some disambiguating information can be propagated from the 'edges of the wallpaper'.

Finding matches between seed points is also done by a species of figural grouping. The procedure employed simply extrapolates from seed points along edge strings except that it can 'jump over gaps' caused by unmatchable edge points. The latter typically arise when left and right image edge points fall outside the allowable orientation limit required for matching (justified by the compatibility constraint and guided by the disparity gradient limit, as in PMF). This tends to happen when the edge string meanders due to local image noise, or closely neighbouring edges cause interference in the locations of Canny edge point locations. When small gaps of this kind are crossed, left image edge points lacking a match due to the gap in the right image do not have a disparity value attached to them.

ORDERING CONSTRAINT

It sometimes happens, particularly in image regions with few edges, that whole strings of edges points are incorrectly matched using the figural continuity procedures outlined above. A check is imposed to discover and remedy such problems using the ordering constraint explicitly and qualitatively. For example, suppose there exist strings A, B, C, D in the left image, each with a strength that is the sum of the strengths of the matches along them (not just length of string). Then if the matches for these strings in the right image violate the ordering constraint (figure 5), then the rule is to kill the weakest strings that result in the ordering constraint being satisfied. The primitives thereby 'released' from matching are considered (along with others also remaining unmatched) in subsequent stages.

Figure 6 Performance on a rock stereogram. (a) Grey level stereo images arranged for cross-eyed fusion. (b) Matched Canny edge points from the left image coded for relative depths with far=light and near=dark. (c) Smooth depth profile obtained from matches in b with the viewpoint from the top left hand side with respect to the other images.

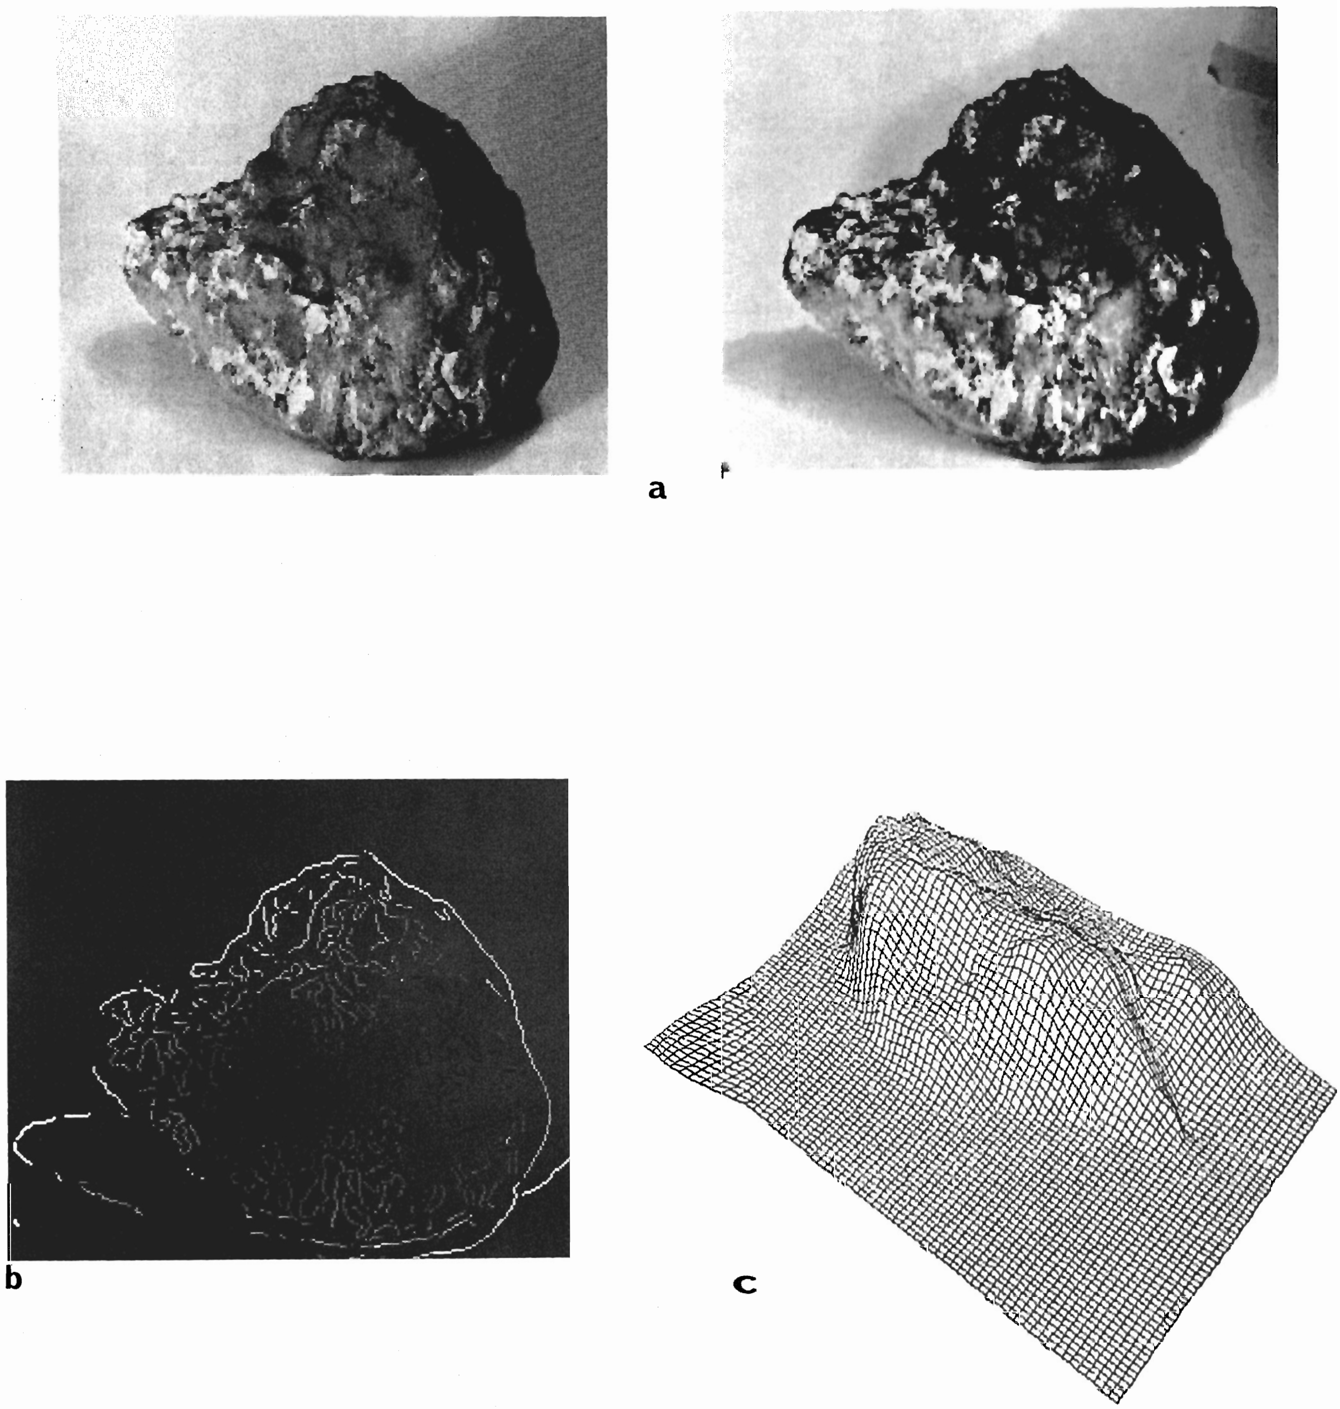

Figure 7 Performance on a cluttered scene of industrial objects. (a) Grey level stereo images arranged for cross-eyed fusion. (b) Canny edges for the left image. (c) Perspective view of depth data recovered by PMF using same program parameters as for the rock in figure 5. (d) Plan view of depth data. (e) Some of the final 3D descriptions fitted to the depth data as recovered by the TINA system (Porrill et al, 1987).

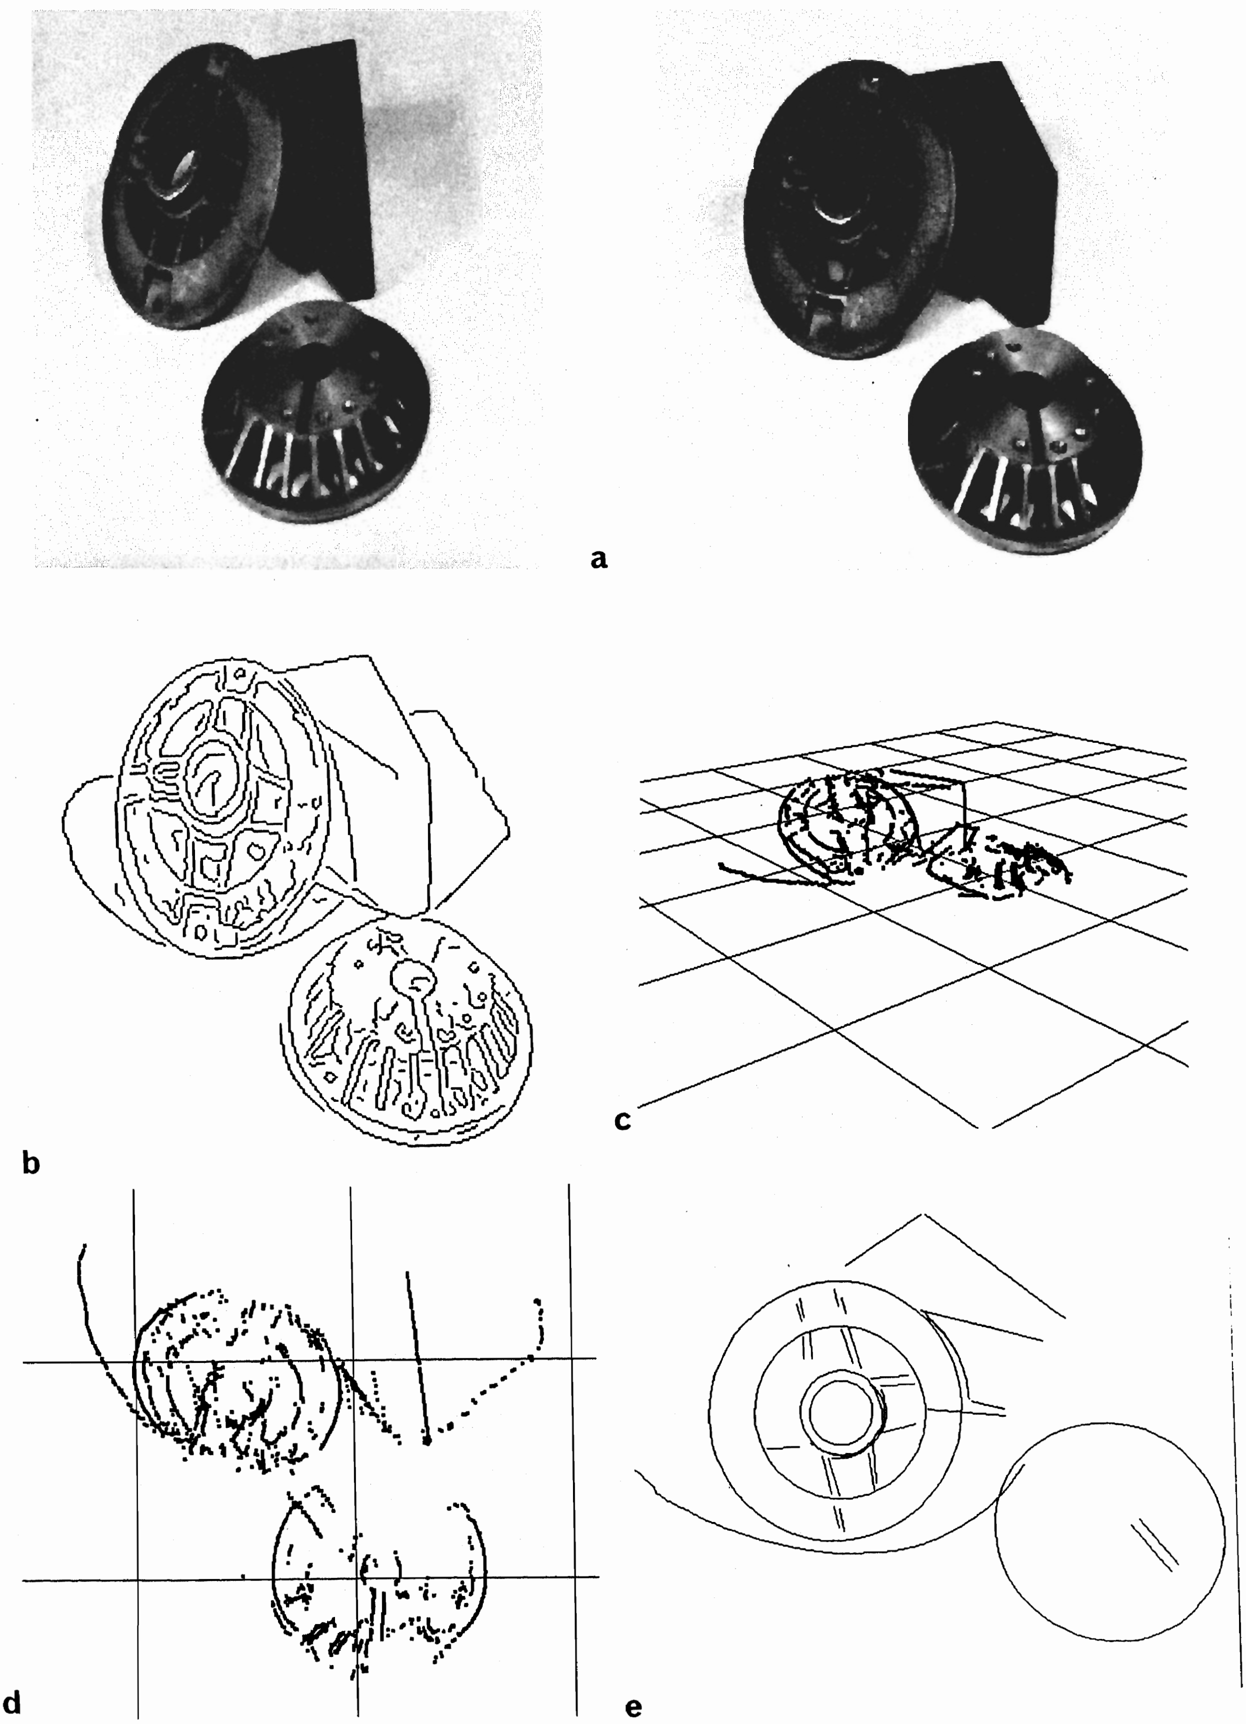

The explicit use of the figural continuity and ordering constraints makes it possible to regard the present implementation as a blend of (1) PMF as it was originally conceived in papers [1] and [2] (i.e. a simple within-DG limit neighbourhood support algorithm); (2) Mayhew and Frisby's (1979) STEREOEDGE algorithm (for explicit use of figural continuity); and (3) the algorithm described by Baker (1982) Baker and Binford (1981) which explicitly exploits the ordering constraint. The original PMF exploited figural continuity and ordering only implicitly (see paper [1]).

COMPUTE DISPARITY RANGE

Some of the remaining ambiguities can be resolved using the ordering constraint to reset the allowable disparity range. That is, matched strings are used to set disparity bounds within which matches must be found for intervening unmatched points if the ordering constraint is not to be violated. This procedure has the advantage of using an appropriate disparity range for each region of the image.

ALLOW CONTRAST-REVERSED MATCHES

Stereo projections do not always preserve the sign of edge contrast. A light-to-dark edge in one image not infrequently projects as a dark-to-light edge in the other image. An example where this happens is at occlusion edges for which the edge is seen against a light background from the vantage point of one eye but against a dark background in the other eye. Thus on the final iteration, matches are allowed between left and right edge points of reversed contrast.

HORIZONTAL EDGES

Horizontal edges present a special problem for stereo matching due to the intractable nature of the matching problem that they pose, viz. there are an infinite number of possible 'solutions' to matching the points comprising horizontal edges. Instead of being assigned a single disparity value, each horizontal edge point in the left image, is assigned a range of possible disparity values determined by the size of the horizontal string in the other image which could be matched with the point in question. These are used in AIVRU's TINATOOL vision system by processes that find best-fit geneotrical descriptions (see paper [11]).

PERFORMANCE EXAMPLES

Figures 6 and 7 show examples of the output achieved by this implementation of the PMF algorithm.

REFERENCES

Baker, H. H. (1982) *Depth from edge and intensity based stereo*. PhD thesis, University of Illinois.

Baker, H. H. and Binford, T. O. (1981) Depth from edge and intensity based stereo. Proceedings of *7th International Joint Conference on Artificial Intelligence*. Los Altos, CA: William Kaufman. Pp 631-636

Mayhew, J.E.W. and Frisby, J.P. (1981) Computation of binocular edges. *Perception* **9** 69-86.

Pollard, S. B. (1985). *Identifying correspondences in binocular stereo.* PhD thesis, University of Sheffield.

Pollard, S. B., Mayhew, J. E. W. & Frisby, J. P. (1985). PMF: a stereo correspondence algorithm using a disparity gradient limit. *Perception* **14** 449-470. (Paper [1] here.)

Pollard, S. B., Porrill, J., Mayhew, J. E. W. & Frisby, J. P. (1985). Disparity gradient, Lipschitz continuity and computing stereo correspondences. In *Proceedings of the Third International Symposium of Robotics Research* (pp. 19-26), Gouvieux, France. Cambridge, MA: MIT Press. (Paper [2] here.)

Shirai, Y. and Nishimoto, Y. (1986) A stereo method using disparity histograms of multi-resolution channels. *Proceedings of Robotics Research: 3rd International Symposium*, pp. 27-32.

Tsai, R. Y. (1986) An efficient and accurate camera calibration technique for 3D machine vision. *Proceedings of the IEEE Computer Vision and Pattern Recognition 86* 364-374.