

# II PMF STEREO ALGORITHM PROJECT

## *Introduction by the Editors*

### I THE GRANT PROPOSAL (Written 1983)

#### 1 OBJECTIVES

The task of a stereo correspondence algorithm is to find matching points in a pair of stereo images. The disparities of matched points can then be measured and used to recover various aspects of the 3D structure of the scene.

The overall goal of the PMF Project is to produce a device capable of delivering in real time (<1 sec) local stereoscopic disparity measurements suitable for the construction of 3D surface descriptions in the 2.5D Sketch Project. PMF (Pollard, Mayhew and Frisby) is the name of the stereo correspondence algorithm which is to form the low level image processing basis of other projects in the Consortium.

The overall goal is to be attained via two research projects. The first aims to explore formally certain mathematical properties of the PMF stereo algorithm with the intention of refining its design. The second is concerned with developing a fast implementation of PMF in special purpose hardware.

#### 2 THE PMF STEREO ALGORITHM

PMF falls into the general category of *neighbourhood support* stereo algorithms, the best known example of which is probably that of Marr and Poggio (1976). The distinctive feature of PMF, however, is to allow potential neighbouring matches to exchange support (mutual facilitation) *if* their relative disparities are not too different. The latter is defined in terms of the disparity between potential matches not exceeding a *disparity gradient limit*. The disparity gradient (DG) between a pair of potential matches is defined in PMF as:

$$DG = \frac{\text{difference in disparities}}{\text{image separation}}$$

A detailed description of PMF is given in Pollard, Mayhew and Frisby (1985)<sup>1</sup>, an abridged version of which is included here as [1] with implementation details of the current version (1990) given in [3]. Papers [1] and [2] discuss how PMF's DG limit can be viewed as enforcing a form of scene-to-image and left-image-to-right-image continuity that can be described as imposing a bound on the degree of 'scene surface jaggedness' to be allowed between selected matches.

In so doing, PMF breaks away from the restrictive notion of surface smoothness implemented by Marr and Poggio's (1976) algorithm. Essentially, they imposed a DG limit of zero by

<sup>1</sup>This paper was submitted for publication at the time of the grant proposal and a good deal of its content was therefore included in the proposal as prior work. References to PMF given in this introduction to the PMF Project are updated versions of equivalent material in the original proposal.

allowing only potential matches with the *same* disparity to exchange support. Restricting facilitation in that way amounts to insisting that surfaces should be 'locally flat and viewed square on' if matching primitives associated with surface markings are to be allowed to exchange support. Yet is evident from simple inspection of tufts of hair, bunches of flowers, etc, that human stereo vision can cope magnificently with scenes that are full of a wide variety of slants and depth discontinuities, scenes comprised of just about anything but (even locally) fronto-parallel planar patches. Using a DG limit of, say, 1.0 (which is the limit reported for binocular fusion in human vision by Burt and Julesz, 1980) is much less restrictive: 1.0 is the maximum DG that can be generated between features on a planar patch with a slant of 84° when viewed at a distance of 65cm with an interocular separation of 6.5cm.

The technical mathematical term for defining surface continuity by reference to a DG limit is Lipschitz continuity (see paper [2]). Moreover, PMF's use of a DG limit can be viewed as a way of parameterising the binocular matching rule of seeking matches which preserve surface 'smoothness' in the Lipschitz sense. If the DG limit is set close to zero then the disambiguating power is great but the range of surfaces that can be dealt with is correspondingly small. If the DG limit is increased to the theoretical limit for opaque surfaces of 2.0 then the range of allowable surfaces is large but disambiguating power is weak because ghost matches then receive and exchange as much support as correct ones. Intermediate values of DG (e.g. 0.5 to 1.5) allow selection of a convenient trade-off point between allowable scene surface jaggedness and disambiguating power because it turns out that most ghost matches produce relatively high DGs (Pollard, 1986).

PMF's use of a DG limit can also be viewed as implementing the ordering, uniqueness and figural continuity constraints used in other stereo algorithms. In addition, PMF uses the DG limit to restrict the number of potential matches entered into its disambiguating support algorithm. It does this by using the DG limit to define an upper bound on allowable orientation differences between left and right edge points when forming potential matches. Papers [1], [2] and [3] provide details on all these issues.

#### 3 CAMERA GEOMETRY

To find matching points requires knowledge of the camera geometry of the stereoscopic imaging device. For the family of camera geometries we shall consider, the correct match of a given point in one image must lie along an 'epipolar line' in the other image. This is illustrated in [1] whose fig 3 shows how left/right pairs of epipolar lines define the locations of all possible matches of points lying along them. In the special case where the principal axes of the cameras are parallel, all epipolar pairs will be horizontal and matching points will be found on corresponding rasters. Indeed, in our implementations of PMF the locations of left and right image eg points are assumed to be rectified to parallel camera

geometry in an image pre-processing stage prior to potential matches being established in order to exploit the simplicity of horizontal epipolars.

For any particular industrial application of the work proposed here, it will either be assumed that the camera geometry is known and fixed, or for applications where it is necessary to have a stereo camera system able to change its convergence, it will be assumed that the development of special purpose techniques suitable for the particular application will be the most economical way of proceeding. Future developments may cast doubt on these assumptions but, for the present at least, the PMF Project does not ask for resources to build a general-purpose stereo camera control system capable of keeping track of its own dynamic geometry by measuring its camera positions and/or by estimating these from analyses of the images themselves.

#### 4 Project A: OPTIMISATION OF PMF

PMF treats the stereo correspondence problem as one of optimising the number of matches subject to the constraint that no pair of matches violates the chosen DG limit. It would in principle be possible to search serially for solutions consistent with this requirement; indeed, one possible way of implementing PMF would be to start at one corner of an image and set up an exhaustive search tree while proceeding through the entire image searching for branches that break the constraint, which would then be eliminated. However, in order to achieve the speed required in a practical application, it is desirable to utilise a parallel algorithm that computes the membership of the optimal set of matches using only local operations.

Considerable investigation of the merits of a DG limit treated as a local constraint has resulted in the adoption of the iterative update scheme set out in [2]. Nevertheless, the updating support equation used in that scheme is not intended to offer a formally satisfactory solution to the global optimisation problem but is simply a way of finding matches that are well supported locally, without violation of a DG limit, and which at the same time can be demonstrated empirically to produce a satisfactory global solution. The support equation used in [2], plus others currently under development that utilise local DG support to select matches, now stands in need of detailed mathematical examination within the framework of constrained optimisation). The objective here is to prove either that this equation (or similar) guarantees a solution of the global DG optimisation problem; or to devise suitable changes that do guarantee this result; or of course to arrive at some principled analysis of inherent limitations in using local DG support if no global optimisation is possible (the latter seems unlikely given the successful demonstrations of PMF to date).

The aim of Project A is to explore this question formally. Dr S A Lloyd, of GEC and who has the requisite mathematical skills and relevant experience in relaxation algorithms, will lead this project under the overall supervision of Dr M McCabe, also of GEC. Lloyd will work in collaboration with Pollard, Mayhew and Frisby of AIVRU.

#### 5 Project B: FAST HARDWARE FOR PMF

It is impossible to give any useful estimates of the potential speed of PMF from the current implementation. The greatest proportion of the code/time involves reading in masses of data

from files and the construction of large structures to represent primitives and their possible matches. But as has already been mentioned, using a DG limit does appear to be open to fast parallel processing techniques, and certainly PMF's iterative update scheme has been designed with this in mind. The fact that less than six iterations are needed to achieve a high measure of consistency for all the examples shown in [1] indicates that a fast implementation of PMF should be attainable and this is the goal of Project A.

This project will be conducted at GEC which is already engaged on various projects for parallel processor architectures and special purpose hardware for signal processing. Specifically, these include the design of a VLSI parallel processor architecture (GRID) and the development of a corresponding parallel programming environment. These facilities and the associated expertise will be exploited in Project A, with PMF's irregular as well as sparse data array posing interesting and fundamental questions with considerable potential relevance to other areas.

Dr J Wijek of GEC will be responsible for ensuring the progress of this work, working in collaboration with Pollard, Mayhew and Frisby on questions relating to PMF's design and performance. Wijek has already been involved with Dr B Buxton (GEC) and Dr H Buxton (Queen Mary College, University of London) in the design of efficient parallel algorithms for the convolutions required by stereo and optical flow computations (implemented on the DAP at Queen Mary College). Close liaison with their work will be maintained.

## II WHAT REALLY HAPPENED?

### 1 Project A: OPTIMISATION OF PMF

This work culminated in the paper by Sheelagh Lloyd [4] which showed that a DG limit could be cast successfully within an optimisation framework.

### 2 Project B: FAST HARDWARE FOR PMF

This project did not develop as planned although the eventual outcome was highly satisfactory.

Immediate changes in project goals were dictated by two main factors. First, Jan Wijek left GEC before completing much work. Secondly, GEC management reviewed their priorities and decided to redeploy resources away from PMF and towards building a pipelined architecture for the Canny edge detector. The reasoning behind this decision was as follows: (a) work on fast hardware for PMF was deemed premature until progress was made on Project A; and (b) as all foreseeable industrial applications of computer vision by GEC would depend on a fast edge detector, effort should first be concentrated in that area. Accordingly, Brendan Ruff was assigned by GEC to develop special purpose hardware for the Canny edge detector and the successful outcome of that work is described in [8]. The upshot of these changes was that developing a version of PMF suitable for the GRID parallel processor did not take place (the GRID project was transferred to other parts of GEC and eventually emerged as a product, MARADE - Marconi Array Demonstrator).

Meanwhile, a collaboration in Sheffield University between its newly appointed Chair of Computer Science, Professor

Doug Lewin<sup>2</sup>, Gordon Manson of that Department, and Chris R Brown in AIVRU, led to a GEC-funded 1-year pilot research project aimed at devising a transputer-based architecture capable in principle of running PMF in under <1 sec. This work produced sufficiently encouraging results (see paper [9] by Brown and Chris Dunford, the latter being an assistant employed by Lewin and Manson) for it to be continued by GEC after the end of the Alvey grant. This continuation took the form of an SERC/ACME-funded GEC/AIVRU collaboration to build MARVIN<sup>3</sup>, a transputer-based fast vision engine for running AIVRU's suite of computer vision programs (called TINATOOL<sup>4</sup>). That device, due in no small measure to Brown's new assistant Mike Rygol who joined AIVRU from INMOS, finally achieved the original Alvey grant objective of PMF in <1 sec, albeit about 2 years after the initial target date. In view of this, a short report on the MARVIN machine is included here as [10], despite that work not being funded under the Alvey grant.

### 3 ADDITIONAL RESEARCH ON THE DISPARITY GRADIENT LIMIT

The optimisation work by Lloyd prompted additional research at GEC into the mathematical properties of the DG limit. Harit Trivedi, a colleague of Lloyd and Margaret McCabe, proved with Lloyd that imposing a DG limit of less than 2 implies that matches preserve the topology of images in the sense of ensuring view-to-view continuity [5].

This result led John Porrill (a postdoctoral research assistant employed on the Alvey grant in AIVRU) to find a simplified proof of the Trivedi-Lloyd theorem and, more importantly, to show that an isotropic DG limit is only one member of a whole family of measures of continuity which impose scene-to-view and view-to-view Lipschitz continuity [2].

### 4 ADDITIONAL RESEARCH AT GEC

Papers [6] and [7] by Trivedi are enclosed which describe methods of estimating stereo and motion parameters. Although not envisaged in the original grant proposal, this work has proved influential and [6] forms the basis for one of AIVRU's current schemes for achieving camera calibration (Thacker and Mayhew, 1990). Trivedi's paper's are included here both for that reason and to illustrate the fact that members of the consortium were able, indeed encouraged, to pursue research topics well beyond those originally conceived in the Alvey grant proposal.

### 5 DEVELOPING THE PMF ALGORITHM

Many developments and evaluations of the PMF algorithm have been conducted by Stephen Pollard in AIVRU over the past 6 years ([1], [2], [3]; see also his PhD thesis: Pollard, 1985).

For example, more thorough empirical investigations of PMF on artificial images [2] confirmed the early claim that

<sup>2</sup> We note with regret that Professor Lewin died before this project was brought to completion.

<sup>3</sup> MARVIN: Multiple ARchitecture for VIsioN

<sup>4</sup> See the introductory overview of the 3D Model-Based Vision Project in Section III for an explanation of this acronym.

enforcing a DG limit well below the theoretical limit for opaque objects of 2 imposes negligible restrictions on the worlds that can be dealt with while at the same time serving admirably to exclude ghost matches, most of which generate DGs above 1.0.

Also, more extended evaluations of PMF have been run on natural images (using Canny edge points as matching primitives instead of the Marr-Hildreth zero crossings used previously). This work continued to demonstrate the power and convenience of using a local neighbourhood support scheme incorporating a DG limit. However, it also demonstrated the desirability of building into PMF more global constraints [3]. Hence procedures exist in the current version of the algorithm which explicitly exploit: (a) figural continuity along strings of edge points (cf. the STEREOEDGE algorithm of Mayhew and Frisby, 1978); and (b) the ordering constraint along epipolars.

Further refinements of the PMF algorithm include: speeded-up processing by restricting initial matching to 'seed points'; allowing matches between primitives of opposite contrast sign when all else fails; and better ways of dealing with the special problems posed by horizontal edges.

Much of this development work on PMF work has been done after the end of the Alvey grant but [3] is included to bring the present account of PMF up to date.

### 6 CONCLUDING REMARKS

The PMF Project played a crucial role in facilitating development and evaluation of the PMF algorithm, itself the backbone of much of the low level image processing work done in AIVRU. The next section of the book (on the 2.5D Sketch Project) describes how the matches generated by PMF were used in the recovery of useful 3D scene geometry for supporting the pick-and-place demonstration that was the culmination of AIVRU's Alvey-supported research (see paper [29]). That demonstration was far from real-time: it took about one hour from capturing a single 256x256 stereo image pair to the robot picking up the object! More recent developments, relying on the MARVIN architecture running much improved code, have brought the self-same demonstration down to 5-10 secs [10].

GEC have benefited from the PMF Project in various ways. The first tangible outcome was their receipt (along with all other participating sites) of a copy of AIVRU's computer vision suite (TINATOOL) at the end of the Alvey project (1987). More recently, continuation via a SERC/ACME grant of the good collaborative relationship built up under Alvey led Bernard Buxton to commission from AIVRU a clone of the MARVIN fast vision engine. That device was delivered to GEC in May 1990, together with a great deal of code for running many component modules of TINATOOL. It is now being used to mount a number of vision based vehicle guidance demonstrators in a collaborative ESPRIT project (VOILA: P2502).

All GEC staff directly working on the PMF Project as it was originally conceived have now left the company. We would particularly like to express our gratitude to Drs McCabe, Lloyd and Trivedi for their assistance in bringing the project to a successful conclusion, albeit not quite the one envisaged at the outset.

## REFERENCES

- Marr, D. and Poggio, T. (1976) Cooperative computation of stereo disparity. *Science* **194** 283-287.
- 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.
- Thacker, N and Mayhew, J E W (1990) Stereo camera calibration from arbitrary images. *AI Vision Research Unit Memo 41*, University of Sheffield.