

4

FINAL MOTION REPORT

CONTRACT NUMBER: N00014-82-K-0464

PROCESSING DYNAMIC IMAGES
FROM CAMERA MOTION

*Massachusetts University
Amherst, MA*

AD-A197 697

DTIC
ELECTE
S JUL 07 1988 D
a D

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

88 05 9

--END

<<ENTER NEXT COMMAND>>

END --

Contents: Overview; Restricted Motion; General Motion and Computation of the Optical Flow Field; Multiresolution Methods; General Motion and Inherent Ambiguities; Processing Approximate Translational Motion for the ALV; Stereoscopic Motion Analysis; Analysis of Constant General Motion; Interpretation; Publications.

1 Introduction and Overview

The motion of points (or other image structures, such as lines or regions) in a two-dimensional image of an evolving three-dimensional environment can be used to find information about that environment, and how it is changing with time. As examples of such information one might want to recover the translational and rotational components of the motion of particular points in the environment, the distance of such points from the image plane of the sensor (the depth), or how the environment is segmented into independent objects.

At the University of Massachusetts, we have developed a number of algorithms for recovery of such information from a sequence of images (frames) of the environment. Our approaches fall generally into two classes. In the first class we follow a two-step paradigm. We first calculate the motion of individual points in the image. This leads to a two-dimensional velocity field in the image called the *optical flow field*. The second step is to use the optical flow field to calculate both the structure of the environment, and the relative motion between frames (its *motion parameters*). In the second class the computation of optical flow and recovery of environmental depth are done concurrently.

Let us initially give an overview of the first class of algorithms mentioned above. Anandan [ANA87c] has developed a method for computing optical flow from a sequence of two images which, in addition to computing a dense flow field (i.e. a velocity vector for each point in the image), uses information in the image to indicate which flow vectors are reliable and which are not. The algorithm has been tested on a wide variety of both synthetic and real image sequences and has been found to be quite robust against the corrupting effects of noise and other image imperfections.



The algorithm of Adiv [ADI85b] uses Anandan's optical flow to robustly compute the sensor motion parameters, the depth of environmental points, object masks for independently moving objects, and the 3D motion parameters of these objects. It is also the only algorithm in the literature that has been demonstrated to compute these quantities in real image sequences containing independently moving objects. Adiv also discusses algorithm-independent ambiguities that arise in the recovery of motion parameters and object masks. This analysis is important in helping to decide in what situations an algorithm can be expected to be effective.

The computation of optical flow and environmental information can also be done concurrently. This is easiest for specific types of motion, for instance when the motion is purely translational. In this case, it is well known that the motion of image points is extremely simple. They move in straight lines away from a single image point known as the Focus of Expansion (FOE). Knowledge of the position of the FOE in the image plane is equivalent to knowing the direction of translation. Lawton [LAW84] has developed an algorithm for computing the position of the FOE. Lawton's FOE-finding algorithm was subsequently improved in the work of Pavlin [PAV86].

When the motion of the camera is purely translational, the depth of environmental points can be found from the position of the FOE and the position of the corresponding image point in two frames of the image sequence. Bharwani, et al. [BHA85,86] developed an algorithm which iteratively refines the computation of environmental depth over a sequence of more than two frames so as to get more and more accurate depths as one progresses through the environment, with a constant computational load between frames. When the camera motion is known to be purely translational, this algorithm is quite

robust.

Unfortunately, in almost all practical imaging situations, it is difficult to ensure that the camera motion is purely translational. A real vehicle moving on a real surface will produce an image sequence in which the image motion will not be due purely to translational motion, but will have an added rotational component. There are two ways to deal with this complex real-world situation. The first way is to eliminate the rotational component of the motion, so that the algorithms which assume purely translational motion can be used. The second way is to use an algorithm which is suitable for general motion (rotation as well as translation). We developed an algorithm (a *Registration Algorithm*) which finds the rotational component of motion, and then removes it from the image sequence so as to produce a new ("registered") image sequence in which the motion is purely translational. We found that the rotational component could not be accurately recovered, and so the registered image sequence still had a small rotational component to its motion. This small rotational component made the determination of the FOE very unstable. This means that the direction of the translational component of the motion is unreliable, and hence so will be the computed depth values. We concluded from this that algorithms which assume purely translational motion will give unreliable results in most practical imaging situations [DUT88].

The only alternative to "registering" an image sequence before applying a purely translational algorithm is to use a general motion algorithm which makes no special assumptions about the motion, i.e. the combination of Anandan and Adiv algorithms. The general motion algorithm was tested on the same sequence of images for which the "registration" had given poor results, and found that the results were significantly improved, an average

error of 8.9% in the depth of a set of points. We are in the process of further testing this algorithm for future use on the Autonomous Land Vehicle.

We now discuss in greater detail the work supported by this contract.

2 Restricted Motion

In the case where the camera motion is purely translational, the finding of the FOE is equivalent to knowing the direction of camera translation. In his Ph.D. Dissertation, Lawton [LAW83,84] formulated an algorithm for finding the position of the FOE which contains two steps: *feature extraction* and *search*. The feature extraction process picks out small image areas which may correspond to distinctive parts of environmental objects. The FOE is then found by a search which determines the image displacement paths along which a measure of feature mismatch is minimized for a set of features. The correct FOE should minimize this error measure and also determine the corresponding image displacement paths for which the extracted features match.

More recently, Pavlin [PAV86] evaluated the Lawton algorithm and found that it could be applied effectively with the analysis of only 8 to 16 image points between frames if the camera is pointed approximately in the direction of camera translation. In addition, the algorithm was speeded up by improving the FOE search algorithm. This was accomplished by computing the error measure for the assumed FOE from a sparser sampling of the visual field (or a more restricted area if constraints on the possible location of the FOE are available). Then, a smooth surface is fit to the error values at those points and the computed minimum of this surface is used to focus the search in the next step of an iterative search process. Both algorithms have been tested on real image sequences with

good results.

When the camera motion is purely translational, knowledge of the FOE and the position of an image point in successive frames of the image sequence allows one to calculate the depth of the corresponding environmental point. Bharwani, et al. [BHA85,86] developed an algorithm which, if the camera motion is purely translational, will compute increasingly more accurate depth information from a sequence of images. This algorithm is intended to be applied after FOE recovery via the Lawton-Pavlin algorithm. The algorithm matches points between frames up to some match resolution, computes a depth range for the environmental point, and then uses this information to predict a smaller search window in future frames, which then can be searched with finer match resolution and consequently more accurate depth. An important characteristic of this algorithm is that the temporal depth refinement can be applied at a constant computational rate and therefore is well-suited for robot navigation.

As we mentioned, however, in a real image sequence, the presence of even small rotational components of the camera motion (of the order of 0.1° to 0.5°) can create havoc on the FOE-finding algorithm. If the FOE cannot be reliably recovered, then the Bharwani (or any other) algorithm cannot be expected to give reliable depth estimates. We discuss this in greater detail in Section 6.

3 General Motion and Computation of the Optical Flow Field

In his recently completed doctoral dissertation [ANA87c] Anandan provides a unified framework for extracting a dense optical flow field from a pair of images, as well as an inte-

grated system which is based on a matching approach (see, also, [ANA85a,b,ANA87a,b]). This framework appears to be sufficiently general to encompass both gradient-based and correlation matching approaches. It consists of a hierarchical scale-based matching scheme using bandpass filters, orientation-dependent confidence measures, and a smoothness constraint for propagating reliable displacements. His integrated system [ANA86] for the extraction of displacement fields uses the minimization of the sum-of-squared-differences (SSD) measure as the local match criterion, and computes confidence measures based on the shape of the SSD surface. It also formulates the smoothness assumption as the minimization of an error functional and overcomes many of the difficult problems that exist in other techniques. The error functional consists of two terms. One, called the *approximation* error, measures how well a given displacement field approximates the local match estimates, while the other, called the *smoothness* error, measures the global spatial variation of a given displacement field. The finite-element method is used to solve the minimization problem. The approach also gives information for extracting occlusion boundaries in some situations.

Anandan has also shown [ANA87d] that the functional minimization problem formulated in his matching technique converges to the minimization problem used in gradient-based techniques (e.g. Glazer's technique discussed later). In particular, by relating an approximation error functional used in his matching approach to the intensity constraints used in the gradient based approaches, he explicitly identifies confidence measures which have thus far been implicitly used in the gradient-based approach. Finally, he suggests ways that algorithms operating on a pair of frames can be developed into multiple-frame algorithms and discusses their relationship to spatio-temporal energy models.

4 Multiresolution Methods

Glazer's recently completed thesis [GLA87c] presents an approach to motion detection using multi-resolution methods in a hierarchical processing architecture. Two motion detection algorithms are developed and analyzed. The hierarchical correlation algorithm utilizes a coarse-to-fine control strategy across the resolution levels and overcomes two disadvantages of single-level correlation: large search areas requiring expensive searches and repetitive image structures which cause incorrect matches. The hierarchical gradient-based algorithm [GLA87a,b], generated over low-pass image pyramids, extends single-level gradient algorithms to the computation of large displacements. Within each level the next refinement of the displacement field is obtained by combining a local intensity constraint and a global smoothness constraint. The mathematical formulation involves the minimization of an error functional consisting of two terms, corresponding to the intensity and the smoothness constraints mentioned above. The minimization problem is solved using the finite-difference approach which leads to a multi-resolution relaxation algorithm. A formal analysis of the hierarchical gradient algorithm is presented, including the basic equations for computing a refined disparity vector, the discrete representations and computations for solving these equations, and a geometric interpretation of the resulting relaxation algorithm. The experimental results show that the two algorithms have comparable accuracy and a cost analysis shows that the hierarchical gradient algorithm is less costly.

5 General Motion and Inherent Ambiguities

In his Ph.D. thesis [ADI85b], Adiv discussed two problems: how to interpret optical flow fields in terms of independently moving rigid objects, and the circumstances under which

such an interpretation is ambiguous. The first problem is important because in real situations the camera may not be the only thing in the environment which is moving, and the second is important because it tells when environmental information can be expected to be unreliable.

Adiv solved the first problem by developing an algorithm which could simultaneously determine the sensor motion parameters, detect (i.e. find object masks) independently moving objects and recover their associated object velocities, and environmental depth [ADI85a]. This algorithm is the only one published which can robustly handle independently moving objects. The algorithm consists of two stages. In the first stage, the flow field is partitioned into connected segments of flow vectors, where each segment is consistent with a rigid motion of a roughly planar surface. Such a segment is assumed to correspond to a part of only one rigid object. This initial organization of the data is utilized in the second stage without the assumption of planar surfaces, and segments are now grouped under the hypothesis that they are induced by a single rigidly moving object and/or by the camera motion. Each hypothesis is tested by searching for environmental motion parameters which are compatible with all the segments in the corresponding group. Once the motion parameters are recovered, the relative environmental depth can be estimated as well. The algorithm was tested on a number of synthetic and real image sequences and found to be able to robustly recover all the environmental parameters.

Adiv then identifies two circumstances under which there are ambiguities in the computation of the environmental parameters [ADI85c]. One such ambiguity occurs when the motion parameters of a sensor or a moving object may be extremely difficult to estimate because there may exist a large set of significantly incorrect solutions which induce flow

fields similar to the correct one. The second occurs when the decomposition of the flow field into sets corresponding to independently moving objects may be ambiguous because two such objects may induce optical flows which are compatible with the same motion parameters. These ambiguity analyses are called "inherent" in the sense that they are algorithm independent. Adiv then discusses the constraints and parameters which can be recovered even in ambiguous situations.

6 Processing Approximate Translational Motion for the ALV

Our previous research in motion analysis led us to attempt to deal with a real application subsystem for the CMU NAVLAB [THO87]. The goal was to detect obstacles in the path of the vehicle at distances beyond the limits of the ERIM range sensor, i.e. at distances beyond 40 feet. Initial results from Bharwani's algorithm implied the possibility of extracting usable depth of obstacles at distances between 40 and 80 feet. By applying an FOE extraction algorithm prior to the depth extraction algorithm, there was the expectation that an effective subsystem could be developed. To accomplish this in actual imaging situations on a moving vehicle turned out to be far more difficult than we expected.

In dynamic imaging situations where the sensor is undergoing primarily translational motion with a relatively small rotational component (i.e. "approximate" translational motion), it might seem likely that translational motion algorithms would be effective in determining depth. Although translational motion is the dominant form of motion and is approximately constant over a long sequence of frames, there usually are local variations due to irregularities in the road surface (bumps, holes, and undulations), as well as minor

rotation of the vehicle as it translates. This is often manifested in changes in the location of the FOE (i.e. effectively it produces a different translational motion), and in rotational motions that must be removed if correct values of depth are to be extracted from the feature displacements. An attempt to correct for these effects via a relatively simple preprocessing "registration" algorithm, without utilizing full analysis of the *general motion problem*, also led to difficulties, even when the rotations were as small as $.1^{\circ}$ to $.5^{\circ}$. These issues and our experimental efforts to deal with what we considered to be the relatively simple problem of approximate translational motion are discussed in [DUT88].

The problems led us recently to apply the Anandan and Adiv algorithms for general motion to the sequences of approximate translational motion, with significantly improved results [DUT88]. The conclusion we draw is that in many real situations general motion analysis must be applied in order to determine depth of points, even when sensor motion is primarily translational with only small amounts of rotation. One obvious hardware solution (at significantly increased cost) is the use of a gyro-stabilized platform or a land navigation system to recover translational and rotational motion so that sensor motion typically will be much closer to the case of pure translational motion. Alternatives to extract motion parameters and depth are outlined in the next section. We will be exploring the utility of these and the general motion algorithm discussed above in the continuation of our work on the Autonomous Land Vehicle.

7 Other Motion Work

7.1 Steroscopic Motion Analysis

By carrying out motion analysis with a pair of cameras - stereoscopic motion - the additional constraints can significantly reduce the complexity of the analysis on a theoretical level. Balasubramanyam and Snyder [BAL87a,b] have developed an algorithm to extract the parameters of *motion in depth*: the single component of translation in depth (i.e. parallel to the line of sight) and the two components of rotation in depth (i.e. rotations that are not around the line of sight). This is achieved by building upon the work of Adiv for segmenting the flow field into rigid independently moving objects [ADI85a], and the formulation of Waxman and Duncan [WAX86], which shows that the ratio of the relative optic flow between a stereo pair of images to the disparity between them is a linear function of the image coordinates. Experimental results are presented for simulated data of general motion of both the sensor and independently moving objects. Work is currently underway to test the effectiveness on real scenes.

7.2 Analysis of Constant General Motion

Another way to introduce additional constraints to the problem of general motion analysis in an effort to achieve practical, robust algorithms is via Shariat's formulation: constant but arbitrary general motion of a rigid object [SHA86]. This leads to a set of difference equations across a sequence of images, relating the positions of a feature in the image plane to the motion parameters of the projected point. The solution obtained is a set of 5th order non-linear polynomial equations in the unknown motion parameters, whose solution requires a Gauss-Newton non-linear least-squares method with carefully designed

initial guess schemes. Pavlin [PAV87] has derived a closed-form solution for the rigid object trajectory by integrating the differential equations describing the motion of a point on the tracked object. The integrated equations are non-linear only in angular velocity, and are linear in all other motion parameters. These equations allow the use of a simple least-square error minimization criterion in an iterative search for the motion parameters.

8 Interpretation

During the life of this project, there have been a variety of efforts that have contributed to the VISIONS system for object recognition and scene interpretation [DRA87,RIS86]. These have included low-, intermediate-, and high-level vision. This is part of our long-range research effort to bring static interpretation together with motion analysis to produce a dynamci interpretation system capable of providing the perception foundation for a mobile robot.

**PUBLICATIONS FOR
FINAL MOTION REPORT**

CONTRACT NUMBER: N00014-82-K-0464

**PROCESSING DYNAMIC IMAGES
FROM CAMERA MOTION**

1987 Publications

- [ANA88a] Anandan, P., "A Unified Perspective on Computational Techniques for the Measurement of Visual Motion", *Proc. of the International Conference on Computer Vision*, London, England, June 1987. Acceptance Rate 13%.
- [ANA87b] Anandan, P., "A Computational Framework and an Algorithm for the Measurement of Visual Motion", COINS Technical Report 87-73, University of Massachusetts at Amherst.
- [ANA87c] Anandan, P., "Measuring Visual Motion From Image Sequences", Ph.D. dissertation and COINS Technical Report 87-21, University of Massachusetts at Amherst, March 1987.
- [ANA87d] Anandan, P., "A Unified Perspective on Computational Techniques for the Measurement of Visual Motion", *Proc. of the DARPA Image Understanding Workshop*, Los Angeles, CA, February 1987, pp. 719-732.
- [BAL87a] Balasubramanyam, Poornima, "Extraction of Motion In Depth: A First Step in Stereoscopic Motion Interpretation", presented at the American Institute of Aeronautics and Astronautics Computers in Aerospace IV Conference, Wakefield, MA, October 7-9, 1987, pp. 277-286.
- [BAL87b] Balasubramanyam, P., "Computation of Motion-In-Depth Parameters Using Stereoscopic Motion Constraints", *Proc. of IEEE Computer Society Workshop on Vision*, Miami, FL, 11/30/87 - 12/2/87.
- [BHA87] Bharwani, S., Arkin, R., Riseman, M., and Hanson, A., "Visual Depth Recovery for Obstacle Avoidance in a Mobile Robot", *Proc. of the IEEE Industrial Electronics Society, Robotics: Trends, Technology and Applications Conference*, Madrid, Spain, September 30-October 1, 1987.
- [BOL87] Boldt, Michael and Weiss, Richard, "Token-Based Extraction of Straight Lines", COINS Technical Report 87-104, University of Massachusetts at Amherst, October 1987.
- [DRA87] Draper B.A., Collins, R.T., Brolio, J., Griffith, J., Hanson, A.R., and Riseman, E.M., "Tools and Experiments in the Knowledge-Directed Interpretation of Road Scenes", *Proc. of the DARPA Image Understanding Workshop*, Los Angeles, CA, February 1987, pp. 178-193. Also COINS TR 87-05, University of Massachusetts at Amherst, January 1987.
- [GIB87a] Giblin, P. and Weiss, R., "Reconstruction of Surfaces from Profiles", *Proc. of the DARPA Image Understanding Workshop*, Los Angeles, CA, February 1987, pp. 900-908. Also COINS TR 87-26, University of Massachusetts at Amherst, March 1987.
- [GIB87b] Giblin, P. and Weiss, R., "Reconstruction of Surfaces from Profiles", *Proc. of the*

International Conference on Computer Vision, London, England, June 1987. Acceptance Rate: less than 25%. Also COINS Technical Report 87-26, University of Massachusetts at Amherst, March 1987.

[GLA87a] Glazer, F., "Hierarchical Gradient-Based Motion Detection", *Proc. of the DARPA Image Understanding Workshop*, Los Angeles, CA, February 1987, pp. 733-748. Also COINS Technical Report 87-77, University of Massachusetts at Amherst.

[GLA87b] Glazer, Frank, "Computation of Optic Flow by Multi-Level Relaxation", COINS Technical Report 87-64, University of Massachusetts at Amherst, July 1987.

[GLA87c] Glazer, Frank, "Hierarchical Motion Detection", Ph.D. dissertation, and COINS Technical Report 87-02, University of Massachusetts at Amherst, February 1987.

[HAN87a] Hanson, Allen and Riseman, Edward, "The VISIONS Image Understanding System", (Chris Brown, Ed.), in *Advances in Computer Vision*, Erlbaum Press, 1987.

[HAN87b] Hanson, A.R., and Riseman, E.M., "Summary of Progress in Image Understanding at the University of Massachusetts", *Proc. of the DARPA Image Understanding Workshop*, Los Angeles, CA, February 1987, pp. 55-64. Also COINS TR 87-20, University of Massachusetts at Amherst, March 1987.

[KAH87] Kahn, P., "Determining Temporal Persistence and Consistent Edge Motion From Natural Images", COINS Technical Report 87-56, University of Massachusetts at Amherst.

[PAV87] Pavlin, Igor, "Motion From a Sequence of Images", COINS Technical Report 87-132, University of Massachusetts at Amherst, February 1988.

[RIS87] Riseman, E., Hanson, A., and Belknap, R., "The Information Fusion Problem: Forming Token Aggregations Across Multiple Representations", submitted to *Computer Vision Graphics and Image Processing*, 1987. Also COINS Technical Report 87-48, University of Massachusetts at Amherst, December 1987.

1986 Publications

[ANA86] Anandan, P., "Computing Optical Flow from Two Frames of an Image Sequence", COINS Technical Report 86-16, University of Massachusetts at Amherst, April 1986.

[BEL86] Belknap, R., Riseman, E., and Hanson, A., "The Information Fusion Problem and Rule-Based Hypotheses Applied to Complex Aggregations of Image Events," *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, June 1986.

[BHA86] Bharwani, S., Riseman, E., and Hanson, A., "Refinement of Environmental Depth Maps Over Multiple Frames," *Proc. of the IEEE Workshop on Motion: Representation and Analysis*, Charleston, SC, May 1986.

[PAV86] Pavlin, I., Riseman, E., and Hanson, A., "Translational Motion Algorithm with Global Feature Constraints", COINS Technical Report 86-58, University of Massachusetts at Amherst, December 1986.

[RIS86a] Riseman, Edward and Hanson, Allen, "The VISIONS Image Understanding System - 1986", COINS Technical Report 86-62, University of Massachusetts at Amherst, December 1986.

[RIS86b] Riseman, Edward M., and Hanson, Allen R., "A Methodology for the Development of General Knowledge-Based Vision Systems", (Carme Torras, Ed.), Alianza Editorial, Spain, 1986. English version of the book to be published by Springer-Verlag.

[WEI86] Weiss, R. and Boldt M., "Geometric Grouping Applied to Straight Lines", *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, Miami Beach, FL, June 1986.

[WIL86] Williams, L. and Anandan, P., "A Coarse to Fine Control Strategy for Stereo and Motion on a Mesh-Connected Computer", *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, June 1986. Also COINS Technical Report 86-19, University of Massachusetts at Amherst, May 1986.

1985 Publications

[ADI85a] Adiv, Gilad, "Determining Three-Dimensional Motion and Structure from Optical Flow Generated by Several Moving Objects", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 4, July 1985, pp. 384-401.

[ADI85b] Adiv, Gilad, "Inherent Ambiguities in Recovering 3-D Motion and Structure From A Noisy Flow Field", *Proc. of IEEE on Computer Vision and Pattern Recognition*, San Francisco, CA, June 19-23, 1985, pp. 70-77.

[ADI85c] Adiv, Gilad, "Interpreting Optical Flow", Ph.D. dissertation and COINS Technical Report 85-35. University of Massachusetts at Amherst, September 1985.

[ANA85a] Anandan, P. and Weiss, R., "Introducing a Smoothness Constraint in a Matching Approach for the Computation of Optic Flow Fields", *Proc. of the IEEE Workshop on Computer Vision*, Bellaire, MI, 1985.

[ANA85b] Anandan, P., "A Review of Motion and Stereopsis Research", COINS Technical Report 85-52, University of Massachusetts at Amherst, December 1985.

[ANA85c] Anandan, P. and Weiss, R., "Introducing a Smoothness Constraint in a Matching Approach for the Computation of Displacement Fields", *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985. Also COINS TR 85-38, University of Massachusetts at Amherst, December 1985.

[BEL85] Belknap, R., Riseman, E., and Hanson, A., "The Information Fusion Problem and Rule-Based Hypotheses Applied to Complex Aggregations of Image Events," *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985.

[BHA85] Bharwani, S., Riseman, E., and Hanson, A., "Refinement of Environmental Depth Maps Over Multiple Frames," *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985.

[HAN85] Hanson, A. and Riseman, E., "Summary of Progress in Image Understanding at the University of Massachusetts", *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985. Also COINS TR 85-48, University of Massachusetts at Amherst, November 1985.

[PAV85] Pavlin, I., Riseman, E., and Hanson, A., "Analysis of an Algorithm for Detection of Translational Motion", *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985.

[WEI85] Weiss, R., Riseman, E., and Hanson, A., "Geometric Grouping of Straight Lines", *Proc. of the DARPA Image Understanding Workshop*, Miami Beach, FL, December 1985.

1984 Publications

[ADI84] Adiv, G., "Recovering 2-D Motion Parameters in Scenes Containing Multiple Moving Objects", *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Washington, D.C., June 1983, pp. 399-400.

[ANA84a] Anandan, P., "Computing Dense Displacement Fields with Confidence Measures in Scenes Containing Occlusion", *Proceedings of SPIE - The International Society for Optical Engineering*, Volume 521, Cambridge, MA, November 1984, pp. 184-194.

[ANA85b] Anandan, P., "Computing Dense Displacement Fields with Confidence Measures in Scenes Containing Occlusion", *Proceedings of the DARPA 1984 Fall Image Understanding Workshop*, New Orleans, LA, October, 1984, pp. 236-247. Also, COINS Technical Report 84-32, University of Massachusetts at Amherst, December 1984.

[LAW84] Lawton, D.T., "Processing Dynamic Image Sequences from a Moving Sensor", Ph.D. dissertation and COINS Technical Report 84-05, University of Massachusetts, February 1984.

[RIE84] Rieger, J.H., and Lawton, D.T., "Processing Differential Image Motion", COINS Technical Report 84-28, University of Massachusetts at Amherst, December 1984.

[RIS84a] Riseman, E.M. and Hanson, A.R., "A Methodology for the Development of General Knowledge -Based Vision Systems", *Proc. of the IEEE Workshop on Principles of*

Knowledge-Based Systems, Denver, Colorado, December 1984.

[RIS84b] Riseman, E.M. and Hanson, A.R., "Image Understanding Research at the University of Massachusetts", *Proc. of the DARPA 1984 Fall Image Understanding Workshop*, New Orleans, Louisiana, October 1984, pp. 1-5.

[WEE84] Weems, Charles C., "Image Processing on a Content Addressable Array Parallel Processor", Ph.D. dissertation and COINS Technical Report 84-14, University of Massachusetts at Amherst, September 1984.

1983 Publications

[GLA83] Glazer, F., Reynolds, G., and Anandan, P., "Scene Matching by Hierarchical Correlation", *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Washington, D.C., June 1983, pp. 432-441.

[HAN83] Hanson, A. and Riseman, E., "A Summary of Image Understanding Research at the University of Massachusetts", *Proc. of the DARPA Image Understanding Workshop*, Arlington, VA, June 1983, pp. 37-42. Also, COINS TR 83-35, University of Massachusetts at Amherst, October 1983.

[LAW83] Lawton, D. T., "Processing Translational Motion Sequences," *Computer Graphics and Image Processing*, Vol. 22, pp. 116-144, 1983.

Non-UMass References

[SHA86] Shariat, H., "The Motion Problem: How to Use More Than Two Frames," Ph. D. Dissertation, Univ. South. Calif., Los Angeles, CA; also, Tech. Rep. IRIS 202, Inst. Rob. and Intell. Syst., Univ. South. Calif., Los Angeles, CA, October 1986.

[THO87] Thorpe, C., Shafer, S., and Kanade, T., "Vision and Navigation for the Carnegie-Mellon NAVLAB," *Proc. DARPA IU Workshop*, Los Angeles, CA, February 1987, pp. 143-152.

[ULL79] Ullman, S., *The Interpretation of Visual Motion*, MIT Press, Cambridge, Mass., 1979.

[WAX86] Waxman, A. N., and Duncan, J. H., "Binocular Image Flows: Steps Toward Stereo-Motion Fusion," *IEEE Trans. on Patt. Anal. and Mach. Intell.*, vol. PAMI-8, November 1986, pp. 715-729.