The Laboratory for Computational Vision

Department of Computer Science
University of British Columbia
Vancouver, B.C. V6T 1W5

D.G. Lowe
A.K. Mackworth
R.J. Woodham

February, 1988

1. Introduction

1.1 History and Growth (1974 -)

Alan Mackworth was appointed in 1974, having just completed his doctoral research under Max Clowes at the University of Sussex. From the outset, research focused on two application domains: interpreting Landsat satellite imagery and interpreting hand-drawn sketch maps. In these domains, the problem of representing and using cartographic and geographic knowledge became central. The Laboratory pioneered applications of the computational vision paradigm to remote sensing.

In Canada and particularly in British Columbia, there is a strong need for trained workers in remote sensing. Accordingly, to strengthen the area at UBC, P.A. Murtha (Forestry and Soil Science), Mackworth, A.L. Farley (Geography) and H.R. Bell (Civil Engineering) jointly proposed an Interdisciplinary Graduate Program in Remote Sensing which was funded by the Universities Council of British Columbia (UCBC) through the UBC Board of Governors as a Program of Excellence in 1977. Under this program three people were hired: Bob Woodham on a joint appointment between the Faculty of Forestry and the Department of Computer Science, with responsibility for the development of the remote sensing program, and two technical support people (one in Murtha's laboratory in Forestry and one in the Laboratory for Computational Vision). These three positions are funded on a continuing basis. Woodham joined the laboratory in 1978, having spent one year as a Research Scientist in the MIT Artificial Intelligence Laboratory after completing his doctoral research there under Berthold Horn in 1977.

It was soon evident that the different research groups at UBC involved in remote sensing required a substantial extension to the digital image analysis computational facility in order to carry out their research programs. Accordingly, Mackworth and Woodham, together with Murtha, W.J. Emery (Oceanography), J.R. Auman, G.G. Fahlman, P. Hickson, T.K. Menon and G.A.H. Walker (Geophysics & Astronomy) and J.L. Paul and W.M. Thurlbeck (Pathology), applied for an NSERC Major Installation Grant entitled, "Digital Image Analysis System Extension". The application was partially funded by NSERC in 1981. Additional funds were supplied by UBC and by MRC.

Bill Havens joined the Laboratory in 1981. Havens completed his doctoral research under Mackworth in 1978. He spent the three years 1978-81 first at Simon Fraser University and then at the University of Wisconsin at Madison where he established and co-directed, with F. Scarpace, the Image Processing Laboratory of the Department of Computer Sciences and the Institute for Environmental Studies. Bill Havens left the Laboratory in 1987 to join Tektronix, Inc. in Portland, Oregon.

David Lowe joined the Laboratory in 1987. Lowe, a UBC B.Sc. Computer Science graduate in 1978, completed his doctoral research at Stanford University under Tom Binford in 1984. From 1984-87, Lowe was an Assistant Professor of Computer Science at the Courant Institute of Mathematical Sciences, New York University.

Mackworth was named the first Director of the Laboratory in 1981. The Director is responsible for overall policy, administration and funding in consultation with other faculty members in the Laboratory. Both Mackworth and Woodham were appointed Fellows of the the Artificial Intelligence and Robotics (AIR) program of the Canadian Institute for Advanced Research (CIAR) in July, 1984. Lowe was appointed an Institute Scholar of the CIAR AIR program in September, 1987. Woodham was named Director of the Laboratory in November, 1984, when Mackworth became Coordinator of the UBC node in the CIAR AIR program. Woodham became Coordinator of the UBC node in the CIAR AIR program, February, 1987.

Since May, 1987, Mackworth and Woodham have continued as Co-Directors of the Laboratory.

Since 1981, research groups in Oceanography, Geophysics & Astronomy, Forestry and Geography have developed computing facilities required for their own remote sensing research programs. Collaboration between the research groups continues, facilitated, in part, by the UBC Remote Sensing Council, Faculty of Graduate Studies.

The experimental component of the Laboratory's research also outgrew the capacity of the equipment resources acquired in 1981. Increasingly, this research required the integration of the computations of early vision, principally numeric in form, with those of knowledge-based vision, principally symbolic in form. Much of our research software is written directly in Lisp or indirectly in knowledge representation systems embedded in Lisp. Woodham and Mackworth, together with U. Ascher, D.G. Kirkpatrick and J.M. Varah (Computer Science) and Murtha, applied for an NSERC Major Equipment Grant entitled, "Lisp Machine for the Laboratory for Computational Vision". The application was funded by NSERC in 1986.

1.2 Equipment and Facilities

The Laboratory for Computational Vision maintains a Digital Image Analysis Facility. The facility supports research in computational vision, remote sensing and other image analysis applications. The computers supported are a DEC Vax 11/780 running Unix 4.3BSD, a Sun 3/160 fileserver with several connected Sun 3/50 and Sun 3/60 workstations, a Sun 2/120FS and a Symbolics 3650 Lisp machine with the Symbolics high resolution colour graphics option and fast convolution hardware from MIT. All these computers are connected via a common Ethernet local area network that is shared with other research computers in the Department of Computer Science. Attached peripherals include an Optronics C-4500 colour film scanner/writer, a Raster Technologies Model One/25 image display workstation, a Matrox MIP512M frame grabber and Hitachi KP321 CCD camera, a Lenco RGB to NTSC video encoder, an Image Resources Videoprint 5200 video hardcopy camera, with both 35mm and SX-70 film systems, and an assortment of other hardware for image input, output and communications. The facility is located in room 202 of the UBC Computer Sciences Building in Department of Computer Science space, physically adjacent to the Computing Centre and its facilities.

The facility provides hardware and software for image input, processing and output on video and a variety of film products. A software library for image manipulation, radiometric and geometric correction, registration, filtering, classification, enhancement and interpretation is supported for Unix systems using a standard image file representation. A standard graphics interface also is supported, extending the basic Unix plot file format. Our standard image file and graphics file software includes interfaces to C and to Franzlisp and support for the Comtal Vision One, the Raster Technologies One/25, telidon, the Optronics C-4500, the Apple Macintosh, magnetic tape, the Imagen laser printer and the Apple LaserWriter. Versions of our Unix standard image software have been provided, under licence, to twenty-one sites, primarily to other universities but including four industrial R&D Laboratories. A variety of software tools for artificial intelligence and computational vision research such as knowledge representation languages also have been developed and supported on Unix systems.

1.3 Users

The facility is primarily available to UBC researchers. Work in computer science, forestry, oceanography, astronomy, geography, physics, electrical engineering, soil science and urban planning has been carried out.

The film scanning and writing capability of the Optronics is unique and has been made widely available to UBC researchers and to researchers from other universities and from private industry. Due to the delicate nature of the Optronics, all film scanning and writing requests are handled by Nedenia Krajci, the Remote Sensing technician. Outside users are billed for film scanning and writing at an hourly rate to recover the cost of technician's time.

External users are screened based on the extent to which the facility contains unique resources essential to their research, the extent to which they can contribute expertise to the facility in, for example, computer networking or programming language development, and the extent to which their research relates to the research activities of the Laboratory for Computational Vision.

The Vax 11/780 system remains our general purpose computer. It is accessible campus-wide via UBCnet and remotely via Datapac.

1.4 Affiliations and Funding

Primary affiliations are with: the UBC Department of Computer Science; the UBC Centre for Integrated Computer Systems Research (CICSR); the UBC Interdisciplinary Graduate Program in Remote Sensing, Faculty of Forestry; and the UBC node in the Artificial Intelligence and Robotics program of the Canadian Institute for Advanced Research.

The operating budget of the Laboratory for fiscal year 87/88 is $156,962 made up as follows:

Source	Amount
NSERC	61,000[a]
University	40,146[b]
CIAR	55,816[c]

[a] From NSERC Infrastructure Grant A0383 to Woodham & Mackworth, (first awarded in 1984-85, renewed for three years in 1987-88)

[b] Technical staff salary support from the UBC Interdisciplinary Graduate Program in Remote Sensing, ($27,036 on-going and $13,110 on-going, subject to annual review, for the duration of Woodham's CIAR Fellowship, currently until June 30, 1989).

[c] Technical & administrative staff salary support from the Canadian Institute for Advanced Research, on-going, subject to annual review.

Additional funding for research supported by the Laboratory is provided by the individual NSERC operating grants of the Faculty members involved.

1.5 Future Plans

The future research plans of the principal faculty members are included in section 3.4 below. Here we describe briefly our plans with respect to future development of the Laboratory.

We plan to continue to restructure our computing equipment into a distributed environment on an Ethernet local area network. This is to improve the access to and productivity of our existing facilities. It is also to facilitate sharing of resources with other research groups and to prepare for the addition of new personnel and facilities both in the Department of Computer Science and in CICSR.

Wolfgang Bibel joined the Department of Computer Science as a Professor and Fellow of the CIAR AIR program in 1987. The Laboratory currently supports his Sun 3/60 on its fileserver. We have a shared interest in AI software, particularly Prolog, and will provide mutual support as Bibel develops his own research facility, based on a Sun 4.

There is one item still to be purchased with funds remaining from our 1986 NSERC Major Equipment Grant. We plan to add Datacube signal processing hardware to the Sun 3/160 fileserver to provide shared access to special purpose image processing hardware.

On-going software development is in Common Lisp, on Sun's and on the Symbolics Lisp Machine, and in Prolog, on Sun's.

2. Staff 1987-88

Co-Directors:	A.K. Mackworth	Computer Science
	R.J. Woodham	Forestry/Computer Science
Faculty:	D.G. Lowe	Computer Science
	A.K. Mackworth	Computer Science
	R.J. Woodham	Forestry/Computer Science
Affiliated Faculty:	P.A. Murtha	Forestry/Soil Science
Guest Faculty:	W. Bibel	Computer Science
	M. Cynader	Ophthalmology (CIAR)
	P.D. Lawrence	Electrical Engineering (CIAR)
	J.A. Mulder	Computer Science (Dalhousie)
	G.W. Neufeld	Computer Science
	R.S. Rosenberg	Computer Science
Other Guests:	D. Brent	Computing Centre
	E. Kuttner	Capilano College
	J. Madden	CIAR Research Council
Post-Doctoral Research Scientists:	E. Eder	Computer Science
	G. Provan	Computer Science
Staff:	T. Lee	CIAR Programmer/Analyst
	N. Krajci	Remote Sensing Technician (1/4 time)
	D. Razzell	Laboratory Manager
	V. McRae	CIAR Secretary
Affiliated Staff:	J. Demco	Computing Centre (CDNNET)
	D. Laplante	Oceanography
	F. Pronk	Computer Science
	R. Sample	Computer Science
	R. Wiart	Forestry
Graduate Students:	J. Brooks	M.Sc.
	E. Fogel	M.Sc.
	A. Gunawan	Ph.D.
	A. Kean	Ph.D.
	M. Majka	Ph.D.

	D. McReynolds	M.Sc.
	F. Mokhtarian	Ph.D.
	J. Mulligan	M.Sc.
	G. Pope	M.Sc.
	M. Romanycia	Ph.D.
	R. Rensink	Ph.D.
	Zhang Y.	M.Sc.
Affiliated Graduate Students:	A. Bansal	Forestry
	N. Dadoun	Computer Science
	A. Gallie	Forestry
	B. Rhone	Computer Science
	U. Runesson	Forestry
	J. Thrower	Forestry
Undergraduate Students:	R. Barman	Computer Engineering
	M. Bolotski	Computer Engineering
	C. Powell	Computer Science

3. Research Activities

3.1 Principal research objective

The objective of the Laboratory for Computational Vision is to make machines see, partly to make them more useful and partly to learn more about the nature of perception itself. The usefulness of machines that see is clear. To us, it is also clear that the ideas of computational vision are the principle scientific path towards making machines see. Moreover, we believe that computational vision is providing a new language for algorithmic theories of perception.

3.2 Areas of specialization

The UBC Laboratory for Computational Vision focuses on applications of computational vision to remote sensing, geographic information systems and robotics. Areas of specialization include: artificial intelligence; knowledge representation; spatial reasoning; model-based vision; constraint propagation; image analysis; and shape representation.

3.3 Major accomplishments to date

Mackworth's early work introduced the use of metric, rather than topological, constraints in the interpretation of line drawings of polyhedral scenes. He advocated the gradient space representation for surface orientation. Gradient space remains one of the basic tools of current research in computational vision.

Many computational vision tasks, as well as others, can best be expressed as constraint satisfaction problems. These problems had typically been solved by backtracking or by poorly understood "filtering" algorithms. Mackworth first described a class of algorithms, known as network consistency algorithms, which provide both a unified framework for analysis and practical techniques to eliminate local inconsistencies in networks of relations.

Woodham's work formally extended Horn's approach to "shape from shading". Woodham also introduced the technique of photometric stereo which has become the first practical realization of methods to determine object shape based on a photometric approach to computational vision.

Woodham has applied his techniques, originally developed for the automatic inspection of industrial parts, to problems in remote sensing. In particular, Woodham has pioneered the use of digital terrain models (DTM's) for both the geometric and the radiometric correction of satellite image data.

Lowe has developed methods for enforcing a viewpoint consistency constraint while matching three-dimensional models to two-dimensional features in an image. New theoretical results in these areas have played a key role in the success of implemented vision systems. One result of this work has been the implementation of a functioning system for solving the industrial bin-of-parts problem, in which a computer can recognize objects from any position even when they are partially occluded behind other objects.

A second area of Lowe's research has been on methods for extracting significant perceptual groupings from low-level image data in order to provide stable image structures for indexing into a database of object models.

3.4 Research programs

The current research programs and future plans of the principal faculty and their students are summarized as follows:

3.4.1 Perceptual Organization and Visual Recognition (D.G. Lowe)

Vision is the most important of the human senses and will eventually play a similarly important role in robotics and a wide range of other tasks requiring computer-based sensing of the environment. My research plan for understanding higher-level vision is to develop capabilities for recognition that can be gradually applied to larger domains of knowledge until they eventually have the same level of performance as human vision. Currently, my work has been at the level of recognizing a few well-defined, rigid objects in a robust manner from arbitrary viewpoints, while making use of only simple line segment data from the image. I am currently improving these methods to handle larger numbers of non-rigid objects matched to arbitrary curves in the image. In the more distant future, I hope to combine more general visual knowledge with an extensive bottom-up analysis that makes use of texture, color, shading and other dimensions of image abstraction. An important component of this task will be to develop methods for learning the appropriate visual knowledge, as this will be the only practical way to acquire the vast amount of visual knowlege necessary for human-level performance.

The following is a more detailed description of several individual research topics:

Perceptual Organization. Human vision is able to spontaneously detect many different types of significant groupings of image elements. For example, people will immediately notice significant instances of parallelism, collinearity, proximity, or symmetry in an otherwise random set of image features. At the theoretical level, my research has been devoted to defining the function that these groupings play in visual recognition and the criteria that are used to determine significance. Based on the principle of detecting structure in the image that is unlikely to have arisen by accident, specific formulas have been derived for measuring the significance of instances of collinearity, proximity of endpoints, and parallelism between straight line segments (Lowe, 1987a). Using these results, we have developed algorithms which can identify these types of groupings in an image, and then match these derived image structures to corresponding structures of a model. Since these image groupings reflect viewpoint-invariant aspects of a three-dimensional scene, they are ideal structures for bridging the gap between the two-dimensional image and the three-dimensional model.

The Viewpoint Consistency Constraint. The viewpoint consistency constraint requires that the locations of all object features in an image must be consistent with projection from a single viewpoint. The application of this constraint is central to the problem of achieving robust recognition, since it allows the spatial information in an image to be compared with prior knowledge of an object's shape to the full degree of available image resolution. In addition, the constraint greatly reduces the size of the search space during model-based matching by allowing a few initial matches to provide tight constraints for the locations of other model features. Unfortunately, while simple to state, this constraint has seldom been effectively applied in model-based computer vision systems.

In previous work (Lowe, 1985d), the author has presented methods for solving for unknown viewpoint and model parameters from a few matches between a three-dimensional object model and a two-dimensional image. In more recent work (Lowe, 1987b), this has been combined with methods for probabilistically evaluating new potential matches to extend and refine an initial viewpoint estimate. This evaluation allows the model-based verification process to proceed without the expense of backtracking or search. The combination of parameter determination along with principled methods for extending an initial match provides for a very reliable final answer as to the presence of an object. Therefore, the remaining aspects of model-based vision are essentially reduced to a problem of minimizing search.

Implementation. The methods of perceptual organization and application of the viewpoint consistency constraint have been combined with many other components of computer vision to produce a functioning model-based vision system, named SCERPO. This system has demonstrated the capability for recognizing known objects from arbitrary viewpoints without any need for depth information. The system is capable of recognizing objects from only partial image data and is quite insensitive to occlusion. The total computation time required to recognize a number of instances of an object in a single image is roughly 3 minutes on a VAX 11/785.

The existing implementation of SCERPO provides a strong base for implementing new techniques in computer vision. There are numerous directions for improving the capabilities and efficiency of the existing system, some of which will be described in this proposal. In addition, the system is a demonstration of the value of several new ideas, such as perceptual organization and viewpoint consistency. We hope to develop these ideas further, particularly in terms of integrating them into a general theory of recognition.

The following subsections outline a number of individual objectives of current research and plans for the near future. Some of these are fairly straightforward extensions to the existing model-based system which will demonstrate the applicability of these methods to a wider class of problems. However, the most ambitious objectives are concerned with the development of a general theory of recognition that would be based on abstracting out the underlying principles from the proven successes of the existing methods. The general theory would extend model-based recognition to many levels of the visual hierarchy and would enable a system to perform recognition in many stages rather than in a single leap.

Incorporating arbitrary curves. Perhaps the most obvious shortcoming of the existing system is its limitation to the use of straight line segments for matching. Much of the theory of matching arbitrary curves has already been developed by the author, such as the ability to solve for viewpoint parameters by minimizing perpendicular errors between curve segments (Lowe, 1987a). However, a substantial amount of work must be undertaken to actually implement these capabilities at all levels of the visual hierarchy. Some of the necessary tasks include low-level curve segmentation and description, perceptual grouping of curve segments, representation of curves in object models, and viewpoint solving based on curve matching.

The payoff from matching with arbitrary curves is much greater than simply extending the generality of the existing system. The reason is that curves are actually much richer features for matching than straight line segments. Often, a match to a single curve will contain enough constraints to solve for all six viewpoint parameters, whereas each straight line segment can only be used to constrain two parameters. By characterizing a curve in terms of zeros, maxima, and minima of curvature, it is possible to produce a description that is largely invariant with respect to viewpoint and can be used to index into a database of models.

Variable model parameters. A second important improvement to the existing system would be to incorporate models with variable sizes, shapes and articulations. The author has previously shown (Lowe, 1980a) how to solve for these variable parameters simultaneously with solving for viewpoint. These methods will allow SCERPO to use models that contain variable sizes and articulations that can contain any combination of translation or rotation parameters. Currently, a modeling system has been implemented that allows for the variable parameters and specific ranges for each parameter, and much work has been completed on enabling the system to solve for these parameters during matching (Goldberg & Lowe, 1987). As a model contains more free parameters, new control methods must be introduced to prevent matching solutions from becoming underconstrained.

Motion tracking. The viewpoint determination and verification components of SCERPO described above can be readily applied to the problem of real-time motion tracking. Limits on the acceleration of an object allow its position to be predicted from frame to frame, thereby minimizing the large search space that must be faced in the general recognition problem. The verification and viewpoint updating can be performed very quickly on an ordinary microprocessor. However, real-time edge detection and segmentation will require special hardware (Lowe, 1987c). We intend to first demonstrate motion tracking off-line with images from a video sequence, and then to subsequently transfer the algorithms to the appropriate hardware.

Wider range of perceptual grouping. The success of the SCERPO system illustrates the value of perceptual organization as an intermediate stage of image analysis. In fact, it is remarkable how useful this organization has proved to be in spite of being limited to the detection of collinearity, parallelism, and connectivity. However, perceptual organization in human vision clearly involves a much wider class of groupings, including curve segmentation, symmetry detection, region formation, and many texture measures. We have already mentioned the desire to implement general curve detection and characterization of viewpoint-invariant properties of curves. This will be supplemented by the implementation of symmetry detection and texture description. All of these forms of grouping will be assigned a level of significance in terms of the degree to which they are unlikely to have arisen by accident from a uniformly distributed background of similar elements.

Hierarchical models. The existing implementation of SCERPO jumps directly from simple image groupings to complete object models during the recognition process. This is feasible only because there are few models under consideration. Human vision clearly performs recognition through a hierarchy ranging from simple components up to full scale object models and scenes.

We plan to extend the object recognition capabilities of SCERPO to incorporate a hierarchy of models. Some of these models would be very simple, such as circles, cylinders or rectangular solids. There would also be more complex partial components of whatever objects were in the current database. The only criteria for such models is that they must be overconstrained in terms of their predictions in the image, so that instances of them can be recognized with confidence. As each simple model is recognized, it would become a feature that could be used to index into the model database to suggest the presence of more complex objects. We believe these methods could greatly reduce the average amount of search required to recognize objects from among a large set of possibilities.

A theory of recognition. A long-term goal is to develop a general theory of recognition that could be applied to all levels of the visual hierarchy. Such a theory is indicated by the commonalities that can be seen in the methods we have used for perceptual organization and the various stages of model-based recognition. In each case, the answer to minimizing search and providing robust performance has been to formulate configurations that are overconstrained with respect to the number of free parameters and can therefore be judged unlikely to have arisen by accident. A careful statistical evaluation of each potential instance of an image feature is used to keep the recognition process focused upon the most reliable data.

While a general theory has not yet been formulated, we can see some of the elements that would play a part. Recognition is performed by formulating models at many levels of the visual hierarchy. Each model is a configuration of features that is unlikely to have arisen by accident, yet has a significant prior probability of being present for particular causal reasons. A Bayesian analysis could show that specific configurations of features are much more likely to have arisen from the presence of a particular object than due to some accidental alignment. While there is a third possibility that the configuration arose due to some non-accidental presence of a different model, this case would be handled separately by also attempting to recognize these competing models. If no such competing models can be identified, then we assume that the features are independent sources of evidence for the presence of the object. The use of these probabilities to learn the most appropriate indexing for an object is described in (Lowe, 1987c).

3.4.2 Knowledge Representations and Algorithms for Computational Vision (A.K. Mackworth)

My research focusses on the theory and applications of computational vision. I have worked primarily on the representation and use of knowledge in vision systems including surface orientation representations, network consistency approximation algorithms for constraint satisfaction problems, schema-based systems and logic-based systems. Applications include diagram understanding, telerobotics and the interpretation of satellite imagery and geographical sketch maps. Recent work focusses on developing a framework to determine descriptive and procedural adequacy for visual knowledge representations, 2-D shape representations, new constraint satisfaction algorithms, complexity analysis of these algorithms, constraint satisfaction algorithms in schema systems for recognition, visual monitoring of an excavator arm, and a model-theoretic logic of depiction.

Under the *cooperative interpretation* paradigm for computational vision and remote sensing, many signal sources with their associated knowledge bases provide constraints on the underlying scene. Such sources include maps and digital terrain models as well as remotely sensed images. One key subtask is to correctly register (or align) maps and images. To this end we have invented, implemented and evaluated the *curvature scale space* representation for planar curves, and used it to match and register a map and a Landsat satellite image automatically (Mokhtarian and Mackworth, 1986a,b). This registration is done manually by most existing systems. In (Mackworth and Mokhtarian, 1987) we show that the 1-D *monotonic property* for scale space can be generalized to curvature scale space, that many properties are invariant under evolution of the curve and that renormalization can handle the so-called 'prickly pear' problem. In a new approach to the shape-from-X problem, using an idea in (Mackworth, 1983b), Katz (M.Sc., 1985) showed how coaxial stereo can generate 3-D environment maps efficiently.

Our constraint satisfaction approach to schema-based vision has led to several new studies and the completion of another Ph.D. thesis (Mulder, 1985). In (Mackworth and Freuder, 1985) we settled a long standing controversy ("Waltz' algorithm *is* linear") and proved, for example, that constraint satisfaction problems can be solved in linear time on trees - a useful result. Network consistency algorithms are *approximation* algorithms enforcing necessary but not always sufficient conditions. In (Mackworth, Mulder and Havens, 1985) we described a new algorithm, hierarchical arc consistency, for finding mutually consistent interpretations in schema networks. We show theoretically that in certain cases it is much faster than arc consistency and give experimental results from Mapsee-3 that support the analysis. In (Mackworth, 1987a) I tie together the recent results in the area and show how all-pervading constraint satisfaction problems are in artificial intelligence. Havens and Mackworth (1987) provide an introduction to Mapsee-2 while Mulder, Mackworth and Havens (1987) provide an overview of the Mapsee project, describing in detail the knowledge structuring techniques developed for Mapsee-1, Mapsee-2 and Mapsee-3.

Under the constraint approach to low-level image interpretation we have shown how to use more realistic image modelling (Woodham, Catanzariti and Mackworth, 1985). Woodham and I have completed a preliminary design for our Sceneflow architecture. Catanzariti, visiting from the University of Naples, and I have studied the use of Gabor functions for edge detection - a crucial stage.

As part of a university-industry applied research project directed by Peter Lawrence, Jane Mulligan, Lawrence and I are designing and implementing a model-based vision system for real-time monitoring of the joint angles on the boom and stick of a Caterpillar excavator.

I previously outlined a proposal for the use of first order logic as a representation language for visual knowledge (Mackworth, 1983a,c). The logic of depiction project (joint with Raymond Reiter) crucially depends on separating out three forms of knowledge: image domain knowledge, scene domain knowledge and the image-scene depiction knowledge. As such it can be used both for graphics and vision. Wong (1986) has built an intelligent graphics system using these ideas. Reiter and Mackworth (1987) describe a formal theory of depiction and interpretation. The theory is illustrated by specifying general knowledge about maps, geographic objects and their depiction relationships in first order logic with equality. An interpretation of an image is defined to be a logical model of the general knowledge and a description of that image. The key advantage of this approach is that it provides a formal framework for analyzing, designing and implementing correct and efficient vision and graphics systems; moreover, it provides a deeper theoretical justification for the roles of constraint satisfaction and satisfiability. We expect the theory to play a major role in our research programs.

In (Mackworth, 1987b) I explore descriptive and procedural adequacy criteria for visual knowledge representation. My research program can be seen in this light - its coherence stems from that. Since almost all the problems we wish to solve are technically hard, we can sacrifice descriptive adequacy to achieve efficiency or we can look for approximation algorithms or other tradeoffs. The work on the logic of depiction enhances both descriptive and procedural adequacy. In (Mackworth, 1987c) those criteria are further developed and used to evaluate some recent work in computational vision.

In the logic of depiction project, Reiter and I are currently evaluating implementation strategies. Target application domains include maps, diagrams and circuit schematics. An important test for the theory is to see how portable the system is across domains. The claim we have made is that a new scene domain requires only a new set of scene, depiction and, possibly, image axioms. The underlying system should be invariant, although the axioms may require hand manipulation to produce an efficent implementation. We are now designing a system for the Sun workstation in consultation with Bibel and Stabler. There are some theoretical issues concerning the tradeoffs between proof-theoretic and model-theoretic approaches to be explored. Apart from providing an experimental testbed, an implementation will be a useful tool in its own right. For example, it will provide a basis for the use of diagrams and circuit schematics in the automation of the process of explanation and diagnosis of complex mechanisms. Another claim to be tested is that the approach can be used both for image analysis and synthesis. A diagram understanding system should be able to interpret a diagram as representing, say, a mechanism and also to generate a diagram from a description of the mechanism using the same underlying knowledge with, perhaps, additional knowledge of formatting conventions.

For the depiction project, I will use my work on constraint satisfaction to provide a module incorporating efficient approximation algorithms for propositional satisfiability. The heuristic importance of these algorithms is that they provide an upper bound on the solution set with computational effort that varies directly and incrementally with the tightness of the bound. I will also work on some open issues here on the degree of parallelism possible in constraint satisfaction algorithms. Although we know that some constraint satisfaction problems are inherently sequential, what remains to be discovered is how large this class is, in theory, and how important this result is, in practice.

Woodham and I plan to continue our work on the Sceneflow architecture. This constraint-based image analysis/synthesis system has a prototype implementation which we are redesigning and planning to test on remote sensing applications. In the curvature scale space project, Mokhtarian and I plan to prove some additional invariant properties of curve evolution that allow it to be used for efficient shape matching. We also plan to explore further the connections between our approach and the 2-D smoothing and curvature approaches of Zucker, Jepson, Lowe, Brady and Asada, Horn and Weldon.

In the visual monitoring project, Mulligan, Lawrence and I are currently at the stage of demonstrating feasibility. We will evaluate the results, apply various engineering criteria and design, implement and evaluate a near real-time system for the Caterpillar excavator. We have already used some of Lowe's work in this project and plan on further collaboration.

3.4.3 Analytic Methods for Image Analysis (R.J. Woodham)

The long-term objective of my research is to identify the constraints and to define the computations that make vision possible, by man or by machine. The strategy followed is to design and to study computer systems that produce descriptions of a 3–D world from 2–D images of that world. For a general purpose vision system, the mapping from signal input to world description is too complex to be treated as a function in a single representation. Many levels of intermediate representation are required. Identifying those levels and establishing the constraints that operate both within and between levels is the fundamental research challenge.

My research demonstrates a problem-driven approach to early vision. For each specific visual task, mathematical equations are derived to model how the 3–D world determines the

2–D image. These equations are based on the physics of image formation and, in general, consider both radiometry and geometry. Recovering information about the 3–D world from 2–D images is the inverse problem. Exact recovery is rarely possible because the inverse problem typically is underconstrained. Several alternatives are possible. In early vision, additional constraint can be obtained from additional images, as in photometric stereo. Additional constraint also can be imposed in the form of a performance criterion (or metric) to determine a unique preferred solution from the space of physically possible solutions.

In general terms, my research contributions have been twofold. First, it is now recognized that a general theory of vision requires careful analysis of the physics of imaging. Not so long ago the prevailing view was that machine vision systems could somehow avoid details of image formation by being "smart". In hindsight, this makes little sense. The physical world imposes constraints. Smartness consists of coming up with methods that exploit those constraints, not of pursuing techniques that ignore them. Second, my work has helped to extend the range of conditions under which vision algorithms can be provably correct. By emphasizing constraints, and the representations necessary to express those constraints, one can extend the effective repertoire of machine vision systems. Photometric stereo and subsequent shape matching using the Extended Gaussian Image (EGI) representation demonstrates this for the robot task of picking mixed parts out of a bin.

The main test-bed for my research is the computer interpretation of remotely sensed data for forest resources management. With the image analysis methods developed, it is possible to decouple direct solar irradiance, diffuse sky irradiance and atmospheric effects from ground cover and topography. When intrinsic surface properties are computed, data acquired at different times and from different locations can be directly compared. The practical result of this research is to extend the range of terrain and imaging conditions that can be handled by automatic image analysis systems.

Research contributions have been in: "shape-from" methods, particularly shape from shading and photometric stereo; the geometric and radiometric correction of satellite multispectral scanner data; shape representations for robotics; and texture perception.

In early vision, the basic scientific question is to determine what can and cannot be computed from image data alone. The intrinsic reflectance of a given object material is specified by its bidirectional reflectance distribution function (BRDF). Given a BRDF, it is possible to derive how that material will appear under any condition of illumination and viewing. Under given conditions of illumination and viewing, the image irradiance equation determines image brightness as a function of surface orientation. The image irradiance equation cannot be inverted locally since image brightness is one measurement and surface orientation has two degrees of freedom. A variety of methods have been developed to compute surface orientation. One method, called photometric stereo, combines multiple images acquired under different conditions of illumination to determine surface orientation locally without requiring a global smoothness constraint (Woodham,1981a). Other methods impose additional constraint in the form of a global smoothness condition (Woodham,1984,1987a). The theme throughout is to develop methods of shape analysis that use all the information available from image data, not just that from a sparse set of discrete features.

As a prerequisite to work in remote sensing and geographic information systems, multiple satellite data sets must be referenced to a common datum, such as that defined by the Universal Transverse Mercator (UTM) projection used in Canadian National Topographic System (NTS) maps. In earlier work, geometric rectification of Landsat multispectral scanner (MSS) imagery to UTM coordinates was achieved by the automatic matching of image features to terrain features obtained from a digital terrain model (DTM) (Little,1982). More recently, these ideas were extended to include orbit refinement strategies and relief displacement correction in the rectification process. This work was the Ph.D. thesis of Frank Wong (Wong,1984) done in collaboration with MacDonald Dettwiler and Associates, Ltd. (MDA) of Richmond, BC. The rectification algorithm developed in Wong's thesis has subsequently been implemented in the Multi-Observational Satellite Image Correction System (MOSAICS) developed by MDA for the Canada Centre for Remote Sensing (CCRS). MOSAICS is the operational system in Canada for the geometric rectification of satellite imagery, including data from the Landsat thematic mapper (TM) and the French SPOT.

An important general task in vision is to separate the effects of object shape and object material both from each other and from effects due to illumination, shadows and viewer position. In remote sensing, image measurements depend on ground cover, elevation, slope and aspect. Methods are required to disambiguate ground cover from topography, independent of the position of the sun and transient atmospheric conditions. Our methods use a digital terrain model (DTM) and, where available, existing forest cover maps. We have used two test sites in British Columbia, one near Shawnigan Lake on Vancouver Island and the other near St. Mary Lake in the Cranbrook timber supply area (TSA). One approach is to use what is known in order to predict what will be seen. The image irradiance equation is a useful tool because it allows one to map knowledge of the world into a common representation, namely the image, to facilitate analysis of what is seen. This "analysis by synthesis" approach was demonstrated for the Shawnigan Lake site (Woodham, Catanzariti & Mackworth,1985). We now also take advantage of detailed forest cover maps obtained from the B.C. Ministry of Forests and Lands. Detailed analysis of the St. Mary Lake site has confirmed our earlier results that surface orientation is the most dominant effect, followed by elevation (Woodham & Lee,1985). More recently, we have implemented numeric models of sky radiance using a uniform tesselation of the sky hemisphere. The calculation of sky irradiance now takes occlusion by adjacent terrain into account. We have shown that skylight and inter-reflection from adjacent terrain too are significant. Finally, we have demonstrated classification map accuracy to operational standards in areas of rugged terrain where traditional methods fail (Woodham & Gray,1987).

Progress has been made on shape representations for use in industrial robotics. A 3–D object can be represented by specifying area as a function of surface orientation, forming its Extended Gaussian Image (EGI). The EGI uniquely represents convex objects, up to translation. J. Little's Ph.D. thesis exploited the EGI representation (Little,1985c). An iterative algorithm for reconstructing a convex polyhedron from its EGI was demonstrated (Little,1983). This algorithm relied on a geometric construction called the "mixed volume". The mixed volume of two convex polyhedra is minimized when the two polyhedra are homothetic (i.e., invariant under translation and scaling). The mixed volume is of basic interest in computational geometry (Little,1985b). We use it as a similarity measure to determine the position and attitude of a known object from a sensed portion of its surface (Little,1985a). Raw surface orientation data is obtained by photometric stereo (Woodham,1981a). This allows a robot to pick mixed parts out of a bin. We determined experimentally that the mixed volume measure of similarity performed better than a direct comparison of area functions. Analysis suggests that the essential difference is one of stability in the measure used to compare shape descriptions. I now view this issue as fundamental and have initiated new studies on stability in measures of similarity and difference between symbolic descriptions of shape (Woodham,1988b).

Recent work has explored both fundamental and applied aspects of texture perception. R. Rensink examined a class of self-similar Gaussian line textures in his M.Sc. thesis (Rensink,1986). Formal analysis has established new relations between the covariance functions and power spectra of stationary self-similar random fields (Rensink,1987). Psychophysical experiment has established conditions of discriminability for the given class of textures. Results are compatible with the hypothesis that human texture perception is mediated by a set of distinct spatial frequency channels. (The initial psychophysical work was done with CIAR Fellow A. Treisman.) Practical aspects of texture-based segmentation of forest scenes were explored in F. Dumoulin's M.F. thesis (Dumoulin,1985). The result was to predict what texture information can, and cannot, be usefully obtained with data from the French SPOT satellite. (This Forestry remote sensing research, done in our Laboratory, is an example of applied work that benefited from direct interaction with our more fundamental studies of texture.)

Most of the projects described above are ongoing. New projects include: work on regularization; sceneflow architecture for computational vision; and colour vision.

According to regularization theory, ill-posed problems can be "solved" by variational principles of a specific type. That is, general theorems exist to guarantee the existence, uniqueness and stability of the solution to suitably "regularized" version of the problem. I have done preliminary work using regularization for edge detection (numerical differentiation), for shape from shading, and for focussing de-focussed images. This work suggests two new ideas.

First, I believe there is a connection to be discovered between optimization by "simulated annealing" and regularization. Everyone knows that it is difficult to find the global extremum of a function that is non-convex. Simulated annealing addresses this problem using a stochastic process to prevent the search from being trapped at local extrema. An alternate approach, like that of regularization, is to smooth the function directly and then to locate the extremum by a standard gradient method. I believe there is a common theory connecting smoothing and the addition of noise. In control theory, for example, it is known pragmatically that high frequency noise, termed "dither", added to a nonlinear system can have the effect of augmenting stability and of quenching undesirable jump phenomena. Zames has shown formally that the stability of a dithered system is related to that of an equivalent smoothed system where the smoothing function is determined by the noise distribution.

Second, smoothing and multi-scale interpretation are closely related concepts. As an example, the regularization parameter used in edge detection can be directly interpreted as the width of the convolution filter used for interpolation prior to differentiation. In this case, there is a direct connection between regularization parameter and degree of smoothing. In general, there is no guarantee that different choices of regularization parameter lead to solutions that differ only in smoothing. This is to be explored in the shape from shading work mentioned above. If a direct connection can be made between regularization parameter and degree of smoothing then the method can be implemented effectively using a multi-grid approach.

Mackworth and I have proposed a sceneflow architecture for computational vision. This is a spreadsheet-like system for image analysis (and synthesis) and is intended as our next general implementation environment. The emphasis is on a flexible control flow that responds automatically to the availability (and change) of various data sources. A paper design has been completed and tested initially using the Flavors object-oriented package in Franz Lisp. A test-bed system is now under development on the Symbolics Lisp machine. Image objects are related by symbolic equations. The interaction of image objects does not depend on how they are implemented internally, allowing for both numeric and symbolic computation. This is our proposed architecture for spatial reasoning and we expect to use it in our work combining remote sensing image analysis and map interpretation. My personal goal in this project is to explore representations at the interface between the numeric and symbolic computations required for vision. In addition to remote sensing, this will also support my work on the computations and representations required for robot bin picking. Ideally, this work also would mesh well with David Lowe's work on the same problem.

A longer term goal is to launch a major effort in colour vision. In my view, an acceptable theory of colour perception must be able to account for the phenomenon of "colour constancy", even when viewing curved surfaces. Also, in my view, a theory of achromatic lightness perception is a prerequisite to a theory of colour perception. These are hard problems but, in some ways, can be seen as logical extensions to my previous work on shape from shading, and, much earlier on, to my work with David Marr on lightness perception. The biological imperative in colour perception is to correctly determine an intrinsic reflectance property of object material, independent of conditions of illumination and viewing. This problem-driven constraint provides the basis for a coherent approach to colour perception.

3.5 Statement on socio-economic relevance

As an academic teaching and research Laboratory, our principal role in the Canadian research community is to provide a research environment for the development and testing of advanced ideas in artificial intelligence, knowledge representation, world modeling, computational vision and graphics. We train the students who become artificial intelligence and robotics researchers in Canadian universities, government agencies and industry.

Our particular areas of expertise are in remote sensing, geographic information systems and robotics. These are areas of immense practical importance to Canada. In particular, our resource-based economy and vast, often inaccessible, land and water bodies make remote sensing the only viable technology for resource management. Canada plays a world leadership role in developing and supplying remote sensing technology. The development of new sensing capabilities in robotics and automation also is crucial to maintaining competitiveness in the global economy. Robot vision is especially important when dealing with non-standard environments typical of forestry, mining and marine applications. At the UBC Laboratory for Computational Vision, we have established a focus of excellence for research and graduate education in these fields.

Our emphasis is on basic research. We hope to contribute to the continual flow of ideas that will give Canada long-term strength in Artificial Intelligence and Robotics.

3.6 Collaboration

The Laboratory is a research group in Computer Science at UBC.

The three CIAR AIR faculty in the Laboratory, Lowe, Mackworth and Woodham, collaborate with P.D. Lawrence in Electrical Engineering, an Associate Fellow of the CIAR AIR program, as well as with several other CIAR AIR Fellows and Associates at other universities across the country. Recently, Max Cynader, also a Fellow of the CIAR AIR program, has joined the UBC Dept. of Ophthalmology, with research facilities located in the Eye Care Centre. We expect strong collaboration in the area of "computational neuroscience".

The Laboratory is affiliated with the UBC Interdisciplinary Graduate Program in Remote Sensing and cooperates with members of the UBC Remote Sensing Council.

The Laboratory collaborates in its remote sensing research with MacDonald, Dettwiler and Associates, Ltd., with the Canada Centre for Remote Sensing and with the B.C. Ministry of Forests and Lands.

The Laboratory collaborates in its vision research with two robotics groups at UBC: D. Cherchas *et al.* in Mechanical Engineering and Lawrence *et al.* in Electrical Engineering. The three groups, and associated faculty, are all part of the recently formed UBC Centre for Integrated Computers Systems Research (CICSR).

Mackworth and Woodham have participated in organizational meetings associated with the recently launched BC Advanced Systems Institute (ASI). Both Mackworth and Woodham are Associate Members of the BC ASI.

4. Publication List

4.1 Journal Articles:

Woodham, R.J. (1988), "What a robot needs to see: theoretical foundations of early vision", *Journal of Mathematical Psychology* (in press).

Lowe, D.G. (1987), "Three-dimensional object recognition from single two-dimensional images", *Artificial Intelligence* (31)355-395.

Lowe, D.G. (1987), "The viewpoint consistency constraint", *International Journal of Computer Vision* (1)57-72.

Lowe, D.G. (1987), "Vision leads robots from the factory", *New Scientist* 115(1577)50-54, (September 10, 1987).

Mackworth, A.K. (1987), "What is the schema for a schema?", (An invited commentary on an article by M.A. Arbib) *The Behavioral and Brain Sciences* (in press).

Woodham, R.J. & M.H. Gray (1987), "Analytic method for radiometric correction of satellite multispectral scanner data", *IEEE Transactions on Geoscience and Remote Sensing* (GE-25)258-271.

Alon A. & W.S. Havens (1986), "Recognizing VLSI circuits from mask artwork", *Revue Internationale de CFAO et d'Infographie* 1(3).

Mokhtarian, F. & A.K. Mackworth (1986), "Authors' reply to comments on Scale-based description and recognition of planar curves and two-dimensional shapes", *IEEE Transactions on Pattern Analysis and Machine Intelligence* (8)675.

Mokhtarian, F. & A.K. Mackworth (1986), "Scale-based description and recognition of planar curves and two-dimensional shapes", *IEEE Transactions on Pattern Analysis and Machine Intelligence* (8)34-43.

Havens, W.S. (1985), "A theory of schema labelling", *Computational Intelligence* (1)127-139.

Hay, J.E. & K.J. Hanson (1985), "Evaluating the solar resource: problems resulting from temporal, spatial and angular variations", *Solar Energy* (34)151-161.

Lowe, D.G. (1985), "Cooperative structuring of information: the representation of reasoning and debate", *International Journal of Man-Machine Studies* (23)97-111.

Lowe, D.G. & T.O. Binford (1985), "The recovery of three-dimensional structure from image curves", *IEEE Transactions on Pattern Analysis and Machine Intelligence* (7)320-326.

Mackworth, A.K. & E.C. Freuder (1985), "The complexity of some polynomial network consistency algorithms for constraint satisfaction problems", *Artificial Intelligence* (25)65-74.

Mackworth, A.K., J.A. Mulder & W.S. Havens (1985), "Hierarchical arc consistency: exploiting structured domains in constraint satisfaction problems", *Computational Intelligence* (1)118-126

Woodham, R.J. (1985), "Trends in quantitative remote sensing", *Forestry Chronicle* (61)363-365.

Woodham, R.J., E. Catanzariti & A.K. Mackworth (1985), "Analysis by synthesis in computational vision with application to remote sensing", *Computational Intelligence* (1)71-79.

Woodham, R.J., & T.K. Lee (1985), "Photometric method for radiometric correction of multispectral scanner data", *Canadian Journal of Remote Sensing* (11)132-161.

Emery, W.J. & M. Ikeda (1984), "A comparison of geometric correction methods for AVHRR imagery", *Canadian Journal of Remote Sensing* (10)46-56.

Ikeda, M., L.A. Mysak & W.J. Emery (1984), "Observation and modelling of satellite-based meanders and eddies off Vancouver Island", *J. Phys. Oceanogr.* (14)3-21.

Ikeda, M., W.J. Emery & L.A. Mysak (1984), "Seasonal variability of the meanders in the California current system off Vancouver Island", *J. Geophys. Res.* (89)3487-3505.

Ikeda, M. & W.J. Emery (1984), "An upwelling event observed in satellite images over the continental shelf off Vancouver Island", *J. Mar. Res.* (42)303-317.

Mackworth, A.K. (1984), "The homunculus as bureaucrat" (A commentary on *The Maltese Cross: A New Simplistic Model for Memory* by D.E. Broadbent), *The Behavioral and Brain Sciences* (1)74.

Raphael, C & J.E. Hay (1984), "An assessment of models which use satellite data to estimate solar irradiance at the Earth's surface", *Journal of Climatology and Applied Meteorology* (23)832-844.

Havens, W.S. (1983), "Recognition mechanisms for schema-based knowledge representations", *International Journal of Computers and Mathematics* (9)185-199

Havens, W.S. & A.K. Mackworth (1983), "Representing knowledge of the visual world", *IEEE Computer* (16)90-96.

Condal, A., G.G. Fahlman & G.A.H. Walker (1981), "The reddening and near-infrared spectrum of NGC 7027" *Publications of the Astronomical Society of the Pacific* (93)191-201.

Woodham, R.J. (1981), "Analysing images of curved surfaces", *Artificial Intelligence* (17)117-140.

Mackworth, A.K. (1980), "Are mediating representations the ghosts in the machine?", *The Behavioral and Brain Sciences* (3)393-394.

Woodham, R.J. (1980), "Photometric method for determining surface orientation from multiple images", *Optical Engineering* 19(1)139-144.

Horn, B.K.P. & R.J. Woodham (1979), "Destriping LANDSAT MSS images by histogram modification", *Computer Graphics and Image Processing* 10(1)69-83.

Dionne, M.S. & A.K. Mackworth (1978), "ANTICS: A system for animating LISP programs", *Computer Graphics and Image Processing* (7)105-119.

Dobrowolski, J.A. & D.G. Lowe (1978), "Optical thin film synthesis program based on the use of Fourier transforms", *Applied Optics* (17)3039-3050.

Starr, D.W. & A.K. Mackworth (1978), "Exploiting spectral, spatial and semantics constraints in the segmentation of Landsat images", *Canadian Journal of Remote Sensing* (4)101-107.

Mackworth, A.K. (1977), "Consistency in networks of relations", *Artificial Intelligence* (8)99-118.

Mackworth, A.K. (1976), "Model-driven interpretation in intelligent vision systems", *Perception* (5)349-370.

Mackworth, A.K. (1973), "Interpreting pictures of polyhedral scenes", *Artificial Intelligence* (4)121-137.

4.2 Books:

Lowe, D.G. (1985), *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, Boston, MA.

4.3 Chapters in Books:

Mackworth, A.K. (1988), "Adequacy criteria for visual knowledge representation", in *Computational Processes in Human Vision: An Interdisciplinary Perspective*, Z. Pylyshyn (ed.), (in press), Ablex Publishing Corp.

Woodham, R.J. (1988), "Stable representation of shape", in *Computational Processes in Human Vision: An Interdisciplinary Perspective*, Z. Pylyshyn (ed.), (in press), Ablex Publishing Corp.

Havens, W.S. & A.K. Mackworth (1987), "Representing and using knowledge of the visual world", in *Knowledge Representation*, N. Cercone & G. McCalla (eds.), pp 429-450, Springer-Verlag.

Lowe, D.G. (1987), "Four steps towards general-purpose robot vision", in *Fourth International Symposium of Robotics Research*, (in press), MIT Press.

Mackworth, A.K. (1987), "Constraint satisfaction", in *Encyclopedia of Artificial Intelligence*, S. Shapiro (ed.), (in press), John Wiley & Sons.

Woodham, R.J. (1987), "Shape analysis", in *Encyclopedia of Artificial Intelligence*, S. Shapiro (ed.), (in press), John Wiley & Sons.

Woodham, R.J. (1984), "Photometric method for determining shape from shading", in *Image Understanding 1984*, S. Ullman & W. Richards (eds.), pp 97-125, Ablex Publishing Corp.

Mackworth, A.K. (1983), "Constraints, descriptions and domain mappings in computational vision", in *Physical and Biological Processing of Images*, O.J. Braddick & A.C. Sleigh (eds.), pp 33-40, Springer-Verlag.

Woodham, R.J. (1983), "Viewer-centred intensity computations", in *Physical and Biological Processing of Images*, O.J. Braddick & A.C. Sleigh (eds.), pp 217-229, Springer-Verlag.

Murtha, P.A. & R.J. Woodham (1983), "Remote-sensing, photo interpretation and photogrammetry", in *Forestry Handbook for British Columbia (4th edition)*, S.B. Watts (ed.), pp 457-500, Forestry Undergraduate Society, UBC.

Mackworth, A.K. (1981), "Consistency in networks of relations", in *Readings in Artificial Intelligence*, B.L. Webber & N.J. Nilsson (eds.), pp 69-78, Tioga Pub. Co.

Woodham, R.J. (1981), "Analysing images of curved surfaces", in *Computer Vision*, J.M. Brady (ed.), pp 117-140, North-Holland.

Woodham, R.J. (1979), "Analyzing curved surfaces using reflectance map techniques", in *Artificial Intelligence: An MIT Perspective*, P.H. Winston & R.H. Brown (eds.), Vol. II, pp 161-182, MIT Press.

Mackworth, A.K. (1978), "Vision research strategy: black magic, metaphors, mechanisms, miniworlds and maps", in *Computer Vision Systems*, A.R. Hanson & E.M. Riseman (eds.), pp 53-61, Academic Press.

Mackworth, A.K. (1977), "How to see a simple world: an exegesis of some computer programs for scene analysis", in *Machine Intelligence 8*, E.W. Elcock & D. Michie (eds.), pp 510-540, John Wiley & Sons.

4.4 Papers in Refereed Conference Proceedings:

Goldberg, R.R. & D.G. Lowe (1987), "Verification of 3-D parametric models in 2-D image data", *Proc. IEEE Computer Society Workshop on Computer Vision*, pp 255-257, Miami Beach, FL.

Rensink, R.A. (1987), "On the visual discrimination of self-similar random textures", *Proc. IEEE Computer Society Workshop on Computer Vision*, pp 240-242, Miami Beach, FL.

Lowe, D.G. (1986), "SYNVIEW: the design of a system for cooperative structuring of information", *Conf. on Computer-Supported Cooperative Work*, pp 376-385, Austin, TX.

Hummel, R.A. & D.G. Lowe (1986), "Computing Gaussian blur", *Proc. 8th Int. Conf. on Pattern Recognition*, pp 910-912, Paris, France.

Havens, W.S. (1986), "Knowledge organization for recognition", *Proc. 1st Australian Artificial Intelligence Congress*, Melbourne, N.S.W., Australia.

Havens, W.S. (1986), "PIPS: A portable image processing system", *IEEE Workshop on Visual Languages*, Dallas, TX.

Little, J.J. (1985), "Extended Gaussian images, mixed volumes, and shape reconstruction", *ACM Symposium on Computational Geometry*, pp 15-23, Baltimore, MD (June 5-7).

Little, J.J. (1985), "Determining object attitude from extended Gaussian images", *Proc. 9th International Joint Conference on Artificial Intelligence*, pp 960-963, Los Angeles, CA (August 20-24).

Lowe, D.G. (1985), "Visual recognition from spatial correspondence and perceptual organization", *Proc. 9th International Joint Conference on Artificial Intelligence*, pp 953-959, Los Angeles, CA (August 20-24).

Woodham, R.J. (1985), "A computational vision approach to remote sensing", *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp 2-12, San Francisco, CA (June 9-13).

Lowe, D.G. (1984), "The representation of debate as a basis for information storage and retrieval", *Proc. National Computer Conference*, pp 595-603, Las Vegas, Nevada.

Mackworth, A.K. & F. Mokhtarian (1984), "Scale-based description of planar curves", *Proc. 5th Canadian Society for Computational Studies of Intelligence Conference*, pp 114-119, London, Ont.

Browse, R.A. (1983), "Computational selection of processing location in vision", *Proc. 5th Annual Conference of the Cognitive Science Society*, Rochester, NY.

Little, J.J. (1983), "An iterative method for reconstructing convex polyhedra from extended Gaussian images", *Proc. American Association for Artificial Intelligence Conference*, pp 247-250, Washington DC.

Lowe, D.G. & T.O. Binford (1983), "Perceptual organization as a basis for visual recognition", *Proc. American Association for Artificial Intelligence Conference*, pp 255-260, Washington DC.

Lowe, D.G. & T.O. Binford (1983), "The perceptual organization of visual images: Segmentation as a basis for recognition", *Proc. DARPA Image Understanding Workshop*, pp 203-209, Arlington, VA.

Mackworth, A.K. (1983), "Recovering the meaning of diagrams and sketches", *Proc. Graphics Interface '83*, pp 313-317, Edmonton, Alta.

Mackworth, A.K. (1983), "On seeing things again", *Proc. 8th International Joint Conference on Artificial Intelligence*, pp 1187-1191, Karlsruhe, West Germany.

Glicksman, J. (1982), "Using multiple information sources in a computational vision system", *Proc. 8th International Joint Conference on Artificial Intelligence*, pp 1078-1080, Karlsruhe, West Germany.

Lowe, D.G. & T.O. Binford (1982), "Segmentation and aggregation: An approach to figure-ground phenomena", *Proc. DARPA Image Understanding Workshop*, pp 168-178, Stanford, CA.

Browse, R.A. (1982), "Interpretation-based interaction between levels of detail", *Proc. 4th Canadian Society for Computational Studies of Intelligence Conference*, pp 27-32, Saskatoon, Sask.

Glicksman, J. (1982), "A schemata-based system for utilizing cooperating knowledge sources in computer vision", *Proc. 4th Canadian Society for Computational Studies of Intelligence Conference*, pp 33-40, Saskatoon, Sask.

Little, J.J (1982), "Automatic registration of Landsat MSS images to digital elevation models", *Proc. IEEE Workshop on Computer Vision: Representation and Control*, pp 178-184, Rindge, NH.

Woodham, R.J. (1982), "Aspects of computational vision", *Proc. 4th Canadian Society for Computational Studies of Intelligence Conference*, pp 237-241, Saskatoon, Sask.

Browse, R.A. (1981), "Relations between schemata-based computational vision and aspects of visual attention", *Proc. 3rd Annual Conference of the Cognitive Science Society*, pp 187-190, Berkeley, CA.

Lowe, D.G. & T.O. Binford (1981), "The interpretation of three-dimen–sional structure from image curves", *Proc. 7th International Joint Conference on Artificial Intelligence*, pp 613-618, Vancouver, BC.

Lowe, D.G. & T.O. Binford (1981), "The interpretation of geometric structure from image boundaries", *Proc. DARPA Image Understanding Workshop*, pp 39-46, Washington, DC.

Mackworth, A.K. & W.S. Havens (1981), "Structuring domain knowledge for visual perception", *Proc. 7th International Joint Conference on Artificial Intelligence*, pp 625-627, Vancouver, BC.

Binford, T.O. R.A. Brooks & D.G. Lowe (1980), "Image understanding via geometric models", *Proc. 5th Int. Conf. on Pattern Recognition*, pp 364-369, Miami, Florida.

Breu, H. & A.K. Mackworth (1980), "Push and pop on pictures: generalizing the augmented transition network formalism to capture the structure and meaning of images", *Proc. 3rd Canadian Society for Computational Studies of Intelligence Conference*, pp 179-187, Victoria, BC.

Browse, R.A. (1980), "Mediation between central and peripheral processes: useful knowledge structures", *Proc. 3rd Canadian Society for Computational Studies of Intelligence Conference*, pp 166-171, Victoria, BC.

Havens, W.S. & A.K. Mackworth (1980), "Schemata-based understanding of hand-drawn sketch maps", *Proc. 3rd Canadian Society for Computational Studies of Intelligence Conference*, pp 172-178, Victoria, BC.

Little, J.J. (1980), "Automatic registration of Landsat images using features selected from digital terrain models", *Proc. 3rd Canadian Society for Computational Studies of Intelligence Conference*, pp 188-195, Victoria, BC.

Lowe, D.G. (1980), "Solving for the parameters of object models from image descriptions", *Proc. DARPA Image Understanding Workshop*, pp 121-127, College Park, MD.

Woodham, R.J. (1980), "Using digital terrain data to model image formation in remote sensing", *Proc. Society of Photo-Optical Instrumentation Engineers* (238)361-369.

Horn, B.K.P. & R.J. Woodham (1979), "Landsat MSS coordinate transformations", *Proc. 5th Machine Processing of Remotely Sensed Data Symposium* pp 59-68, West Lafayette, Ind.

Woodham, R.J. (1979), "Relating properties of surface curvature to image intensity", *Proc. 6th International Joint Conference on Artificial Intelligence* pp 971-977, Tokyo, Japan.

Browse, R.A. (1978), "Knowledge identification and metaphor", *Proc. 2nd Canadian Society for Computational Studies of Intelligence Conference*, pp 48-54, Toronto, Ont.

Catanzariti, E. & A.K. Mackworth (1978), "Forests and pyramids: using image hierarchies to understand Landsat images", *Proc. 5th Canadian Symposium on Remote Sensing*, pp 284-291, Victoria, BC.

Havens, W.S. (1978), "A procedural model of recognition for machine perception", *Proc. 2nd Canadian Society for Computational Studies of Intelligence Conference*, pp 254-262, Toronto, Ont.

Mackworth, A.K. (1978), "Computer-based image understanding at UBC", *Proc. Canadian Information Processing Society '78*, pp 422-424, Edmonton, Alta.

Mulder, J.A. & A.K. Mackworth (1978), "Using Multi-level semantics to understand sketches of houses and other polyhedral objects", *Proc. 2nd Canadian Society for Computational Studies of Intelligence Conference*, pp 244-253, Toronto, Ont.

Woodham, R.J. (1978), "Photometric stereo: a reflectance map technique for determining object relief from image intensity", *Proc. 2nd Canadian Society for Computational Studies of Intelligence Conference*, pp 117-124, Toronto, Ont.

Woodham, R.J. (1978), "Photometric stereo: a reflectance map technique for determining surface orientation from image intensity", *Proc. Society of Photo-Optical Instrumentation Engineers* (155)136-143.

Havens, W.S. (1977), "A procedural model of recognition", *Proc. 5th International Joint Conference on Artificial Intelligence* pp 264, Cambridge, MA.

Mackworth, A.K. (1977), "On reading sketch maps", *Proc. 5th International Joint Conference on Artificial Intelligence* pp 597-606, Cambridge, MA.

Woodham, R.J. (1977), "A cooperative algorithm for determining surface orientation from a single view", *Proc. 5th International Joint Conference on Artificial Intelligence* pp 635-641, Cambridge, MA.

Havens, W.S. (1976), "Can frames solve the chicken and egg problem?", *Proc. 1st Canadian Society for Computational Studies of Intelligence Conference*, pp 232-242, Vancouver, BC.

Mackworth, A.K. (1976), "Making maps make sense", *Proc. 1st Canadian Society for Computational Studies of Intelligence Conference*, pp 42-51, Vancouver, BC.

4.5 Department of Computer Science Technical Reports:

Kuttner, E. (1987), "A schema and constraint based representation to natural language understanding", TR-87-2.

Mackworth, A.K. (1987), "Adequacy criteria for visual knowledge representation", TR-87-4.

Mackworth, A.K. (1987), "Update on computational vision: shape representation, object recognition and constraint satisfaction", TR-87-29.

Mulder, J.A., A.K. Mackworth & W.S. Havens (1987), "Knowledge structuring and constraint satisfaction: the mapsee approach", TR-87-21.

Reiter, R. & A.K. Mackworth (1987), "The logic of depiction", TR-87-24.

Romanycia, M.H.J. (1987), "The design and control of visual routines for the computation of simple geometric properties and relations" TR-87-34.

Woodham, R.J. (1987), "Stable representation of shape", TR-87-5.

Woodham, R.J. & M.H. Gray (1987), "Analytic method for radiometric correction of satellite multispectral scanner data", TR-87-1.

Rensink, R. (1986), "On the visual discrimination of self-similar random textures", TR-86-16.

Woodham, R.J. (1986), "Shape analysis", TR-86-12.

Alon, A. & W.S. Havens (1985), "Recognizing VLSI circuits from mask artwork", TR-85-1.

Havens, W.S. (1985), "A portable image processing system for computer vision", TR-85-9

Katz, I. (1985), "Coaxial stereo and scale-based matching", TR-85-13.

Little, J. (1985), "Recovering shape and determining attitude from extended Gaussian images", TR-85-2.

Mackworth, A.K. (1985), "Constraint satisfaction", TR-85-15.

Mackworth, A.K., J. Mulder & W.S. Havens (1985), "Hierarchical arc consistency: exploiting structured domains in constraint satisfaction problems", TR-85-7.

Mulder, J. (1985), "Using discrimination graphs to represent visual knowledge", TR-85-14.

Havens, W. (1984), "A theory of schema labelling", TR-84-16.
Mackworth, A. & F. Mokhtarian (1984), "Scale-based descriptions of planar curves", TR-84-1.
Mokhtarian, F. & A.K. Mackworth (1984), "Scale-based description and recognition of planar curves and two-dimensional shapes", TR-84-15.
Wong, F. (1984), "A unified approach to the geometric rectification of remotely sensed imagery", TR-84-6.
Woodham, R.J. (1984), "Photometric method for determining shape from shading", TR-84-10.
Woodham, R.J. & T.K. Lee (1984), "Photometric method for radiometric correction of multispectral scanner data", TR-84-14.
Havens, W. (1982), "Standard image files", TR-82-10.
Mackworth, A.K. (1982), "The complexity of some polynomial network consistency algorithms for constraint satisfaction problems", TR-82-6.
Glicksman, J. (1983), "A cooperative scheme for image understanding using multiple sources of information", TR-82-13.
Browse, R. (1982), "Knowledge-based visual interpretation using declarative schemata", TR-82-12.
Little, J. (1980), "Automatic rectification of Landsat images using features derived from digital terrain models", TR-80-10.
Catanzariti, E. & A.K. Mackworth (1978), "Forests and pyramids: using image hierarchies to understand Landsat images", TR-78-2.
Havens, W.S. (1978), "A procedural model of recognition for machine perception", TR-78-3.
Starr, D. & A.K. Mackworth (1978), "Exploiting spectral, spatial and semantic constraints in the segmentation of Landsat images", TR-78-1.
Mackworth, A.K. (1977), "On reading sketch maps", TR-77-2.
Mackworth A.K. (1976), "Model-driven interpretation in intelligent vision systems", TR-76-2.
Starr, D.W. & A.K. Mackworth (1976), "Interpretation-directed segmentation of ERTS images", TR-76-3.
Mackworth A.K. (1975), "Consistency in networks of relations", TR-75-3.
Mackworth A.K. (1975), "How to see a simple world", TR-75-4.

4.6 Graduate Theses:

Bult, T. (M.Sc. 1987), "Schema labelling applied to hand-printed Chinese character recognition" (supervisor W.S. Havens).

Crudge, S. (M.Sc. 1987), "Quantification of rill erosion using field measurements and remote sensing techniques" (supervisor H. Schreier).

Donahue, J. (M.A.Sc. 1987), "The extraction of landslides in a satellite image using a digital elevation model" (supervisor P.A. Murtha).

Collins, M. (M.Sc. 1987), "The estimation of pack-ice motion in digital satellite imagery by matched filtering" (supervisor W.J. Emery).

Romanycia, M. (M.Sc. 1987), "The design and control of visual routines for the computation of simple geometric properties and relations" (supervisor A.K. Mackworth).

Thomas, A. (Ph.D. 1987), "Relationship between near-surface plankton distribution, hydrography and satellite measured sea-surface thermo patterns" (supervisor W.J. Emery).

Vachon, P. (Ph.D. 1987), "Synthetic aperture radar imaging of the ocean surface: theoretical considerations, and experiments with simulated and actual SAR imagery" (supervisor W.J. Emery).

Gray, M. (M.Sc. 1986), "Radiometric correction of satellite imagery for topographic and atmospheric effects" (supervisor R.J. Woodham).

Kuttner, E. (M.Sc. 1986), "A schema and constraint based representation to natural language understanding" (supervisor W.S. Havens).

Rensink, R. (M.Sc. 1986), "On the visual discrimination of self-similar random textures" (supervisor R.J. Woodham).

Shin, H-Y. (Ph.D. 1986), "Satellite derived sea surface temperature: a physical approach" (supervisor W.J. Emery).

Wong, G. (M.Sc. 1986), "Depiction and domains in visual knowledge representation" (supervisor A.K. Mackworth).

Benchimol, N. (M.Sc. 1985), "The mesoscale variability of insolation over the lower Fraser Valley resolved by geostationary satellite data" (supervisor J.E. Hay).

Dumoulin, F. (M.F. 1985), "Using texture energy measures for the segmentation of forest scenes" (supervisor R.J. Woodham).

Heiss, D. (M.Sc. 1985), "Calibrating the photographic reproduction of colour digital images" (supervisor R.J. Woodham).

Mulder, J. (Ph.D. 1985), "Using discrimination graphs to represent visual knowledge" (supervisor A.K. Mackworth).

Katz, I. (M.Sc. 1985), "Coaxial stereo and scale-based matching" (supervisor A.K. Mackworth).

Little, J. (Ph.D. 1985), "Recovering shape and determining attitude from extended Gaussian images" (supervisor R.J. Woodham).

Archibald, C. (M.Sc. 1984), "Using range data for object recognition in a robotics environment" (supervisor A.K. Mackworth).

Girard, J. (M.Sc. 1984), "A recursively controlled production system" (supervisor W.S. Havens).

Jankowski, R. (M.A.Sc. 1984), "Displaying human visual field data as shaded surfaces" (supervisors G.F. Schrack & R.J. Woodham).

Mokhtarian, F. (M.Sc. 1984), "Scale-based description and recognition of planar curves" (supervisor A.K. Mackworth).

Stewart, K.E. (M.Sc. 1984), "The angular distribution of diffuse solar radiation over the sky hemisphere" (supervisor J.E. Hay).

Wanless, N. (M.Sc. 1984), "Satellite based estimates of solar energy availability at the Earth's surface. Impact of changes in the configuration of the satellite data" (supervisor J.E. Hay).

Wong, F. (Ph.D. 1984), "A unified approach to the geometric rectification of remotely sensed imagery" (supervisor R.J. Woodham).

Lee, T. (M.Sc. 1983), "Modeling and estimating atmospheric effects in Landsat imagery" (supervisor R.J. Woodham).

Sims, R.A. (Ph.D. 1983), "Ground truth and large-scale 70mm aerial photographs in the study of reindeer winter rangeland, Tuktoyaktuk Pennisula area, N.W.T." (supervisor P.A. Murtha)

Browse, R.A. (Ph.D. 1982), "Knowledge-based visual interpretation using declarative schemata" (supervisor A.K. Mackworth).

Glicksman, J. (Ph.D. 1982), "A cooperative scheme for image understanding using multiple sources of information" (supervisor A.K. Mackworth).

Majka, M. (M.Sc. 1982), "Reasoning about spatial relationships in the primal sketch" (supervisor R.J. Woodham).

Palmer, M. (M.Sc. 1982), "Using surface models to alter the geometry of real images" (supervisor R.J. Woodham).

Palmer, P (M.Sc. 1982), "Colour manipulation of digital images" (supervisor R.J. Woodham).

Raphael, C. (M.Sc. 1982), "Models for estimating solar irradiance at the Earth's surface from satellite data: An initial assessment" (supervisor J.E. Hay).

Stoch, L. (M.Sc. 1981), "Interactive polygon filling on a raster graphics display" (supervisor R.J. Woodham).

Breu, H. (M.Sc. 1980), "Slant from texture: computational methods for recovering surface slant from images of textured scenes" (supervisor A.K. Mackworth).

Koomen, J.A.G.M. (M.Sc. 1980), "The interlisp virtual machine: a study of its design and implementation as multilisp" (supervisor A.K. Mackworth).

Leckie, D.G. (Ph.D. 1980), "Development of a night-time cooling model for remote sensing thermal inertia mapping" (supervisor P.A. Murtha).

Little, J. (M.Sc. 1980), "Automatic rectification of Landsat imagery using features derived from digital terrain models" (supervisor R.J. Woodham).

Fine, G. (M.Sc. 1979), "Design of an intelligent LISP CAI tutor" (supervisor A.K. Mackworth).

Pollack, K.A. (M.Sc. 1979), "LOGO: an approach to computer-based learning" (supervisor A.K. Mackworth).

Mulder, J.A. (M.Sc. 1979), "Representation and control in a program that understands line sketches of houses" (supervisor A.K. Mackworth).

Barenholtz, J. (M.Sc. 1978), "Towards a new animation technology" (supervisor A.K. Mackworth).

Havens, W.S. (Ph.D. 1978), "A procedural model of recognition for machine perception" (supervisor A.K. Mackworth).

Catanzariti, E. (M.Sc. 1977), "Using image hierarchies to interpret Landsat data" (supervisor A.K. Mackworth).

Starr, D.W. (M.Sc. 1976), "Automatic interpretation of Landsat images using context-sensitive region merging" (supervisor A.K. Mackworth).

Dionne, M.S. (M.Sc. 1975), "ANTICS - a system for animating LISP programs" (supervisor A.K. Mackworth).