1985 IEEE Workshop on Languages for Automation

Cognitive Aspects in Information Processing

PROCEEDINGS

1985 IEEE WORKSHOP ON LA[NGUAGES] FOR AUTOMATION

COGNITIVE ASPECTS IN INFORMATION PROCESSING

UNIVERSITAT DE PALMA
PALMA DE MALLORCA, SPAIN
JUNE 28-29, 1985

LANGUAGES LANGUAGES LANGUAGES LANG

NG PROGRAMMING PROGRAM

GN DESIGN DESIGN DESIGN

ES, HIGH-LEVEL COMMUNICATION LANG

AGES ROBOTIC CONTROL LANGUAGES ROBOT

ISBN 0-8186-0619-3
IEEE CATALOG NO. 85CH2141-0
LIBRARY OF CONGRESS NO. 85-60199
IEEE COMPUTER SOCIETY ORDER NO. 619

IEEE COMPUTER SOCIETY

COMPUTER SOCIETY PRESS

1884 1984

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

The papers appearing in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and are published as presented and without change, in the interests of timely dissemination. Their inclusion in this publication does not necessarily constitute endorsement by the editors, IEEE Computer Society Press, or the Institute of Electrical and Electronics Engineers, Inc.

Published by IEEE Computer Society Press
1109 Spring Street
Suite 300
Silver Spring, MD 20910

ISBN 0-8186-0619-3 (Paper)
ISBN 0-8186-4619-5 (Microfiche)
ISBN 0-8186-8619-7 (Casebound)
IEEE Catalog Number 85CH2141-0
Library of Congress Number 85-60199
IEEE Computer Society Order Number 619

Order from: IEEE Computer Society
Post Office Box 80452
Worldway Postal Center
Los Angeles, CA 90080

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

Preface

The annual **IEEE Workshop on Languages for Automation** covers issues related to the overall underpinning role of languages for advanced forms of automation and robotics. The theme of LFA 1985, "Cognitive Aspects in Information Processing," focuses on research having to do with filtering, storing and attending relevant information, and topics such as models and algorithms for decision making, behavioral signature extraction, shift of attention among/within behavioral modalities, filtering, and zooming and focusing operations.

This year the workshop was held at the Universitat de Palma, Mallorca, Spain, on June 28-29, 1985. Next year's **IEEE Workshop on Languages for Automation** will be hosted by the Institute of System Science of the University of Singapore in August 1986.

On behalf of the program committee, we would like to thank all the participants and contributors who have made this workshop a great success. We also wish to thank Dr. Nadal Batle, Rectorat Universitat de Palma, for the gracious hospitality extended to us by the Universitat de Palma, and to Professor T. Riera who so diligently managed all local arrangements for the **1985 IEEE Workshop on Languages for Automation.**

P.A. Ligomenides
S.K. Chang
Program Cochairmen

Conference Committee

Workshop Honorary Chairman

L.A. Zadeh, *University of California, Berkeley, USA*

Workshop General Chairman

H.J. Zimmermann, *Tech. U. Aachen, Germany*

Program Committee Chairman

P.A. Ligomenides, *University of Maryland, USA*

Program Committee Cochairman

S.K. Chang, *Illinois Institute of Technology, USA*

Keynote Speaker

A. Rosenfeld, *University of Maryland, USA*

Local Arrangements

T. Riera, *Universitat Politecnica de Catalunya, Spain*

Program Committee Members

J.C. Bezdek, *University of South Carolina, USA*
K.S. Fu, *Purdue University, USA*
J. Gutierrez, *Universidad Politecnica de Madrid, Spain*
S.M. Kosslyn, *Harvard University, USA*
T. Riera, *Universitat Politecnica de Catalunya, Spain*
A. Rosenfeld, *University of Maryland, USA*
R.R. Yager, *Iona College, USA*

Table of Contents

ORIGINAL PAGE IS
OF POOR QUALITY

*Not received in time for publication

Knowledge System Models
(Chairman: P.A. Ligomenides, University of Maryland)

Expert Systems

(This session is dedicated to the memory of Dr. King-sun Fu, Goss Distinguished Professor of Engineering and Professor of Electrical Engineering at Purdue University in West Lafayette, Indiana. Dr. Fu, one of the organizers of this conference, died April 29, 1985.)

Languages and Simulation
(Chairman: H.J. Zimmermann, Tech. U. Aachen, West Germany)

Robot Vision
(Chairman: S.-S. Chen, National Science Foundation)

Keynote Address

ROBOT VISION AND SENSORY DATA ANALYSIS

Azriel Rosenfeld
Center for Automation Research
University of Maryland
College Park, MD 20742 U.S.A.

Computer vision systems for robotic and other applications are rapidly becoming more powerful and more widely applicable. This paper defines low-, intermediate-, and high-level vision systems and indicates where progress is needed, especially in the third of these areas.

INTRODUCTION

Computer vision, originally a branch of pattern recognition, is now over 30 years old. It has produced many practical accomplishments in areas such as document processing (character recognition), medicine (e.g., blood cell counting), remote sensing, and industrial automation (inspection, robot vision, etc.).

The oldest computer vision systems dealt with (approximately) two- dimensional input data. Examples of such inputs include documents, cross-sections of microscopic specimens, high-altitude views of the earth's surface, and circuit boards or chips. The analysis of two-dimensional scenes is now generally known as "low-level" vision.

During the past 15 to 20 years, a variety of techniques have been developed for dealing with images of three-dimensional scenes, such as arise in most robot vision applications. In order to analyze such scenes, it is generally necessary to solve two difficult problems: (1) to determine the three-dimensional shapes of the visible surfaces in the scene on the basis of the information available in the image(s); (2) to recognize objects in the scene when their spatial orientations are not known and they are only partially visible. The techniques developed to solve these problems and thus to analyze 3D scenes are generally known as "intermediate-level" vision techniques.

To recognize objects in a scene (whether two- or three-dimensional), a vision system must incorporate knowledge about the class of scenes to be analyzed and the types of objects that occur in them. In most existing vision systems, this knowledge is implicit: it is "embedded" in the analysis procedures, which are designed to perform well when applied to scenes of the given class. Recently there has been increased interest in developing "expert" vision systems in which the knowledge about the scene, as well as the rules for using this knowledge to analyze the image(s), are represented explicitly. We shall call such systems "high-level" vision systems.

In the next two sections of this paper, block diagrams for generic low- and intermediate- level vision systems are presented and briefly discussed. The last section presents some general comments on high-level vision systems. For simplicity, we shall assume throughout that the input is a single optical image; but the discussion can easily be extended to images obtained from other sensors (e.g., range images), to multiple images (e.g., stereopairs), or to time sequences of images.

LOW-LEVEL VISION

A block diagram of a generic low-level vision system for the analysis of two-dimensional scenes is presented in Figure 1. The main stages in such a system are as follows:

a) Features (such as edges or curves) are extracted from the image, or it is segmented into regions of various types (light vs. dark, smooth vs. textured, etc.). The result of this stage is a "feature map" or "symbolic image" in which pixel values represent feature or region labels rather than gray levels.

b) Properties of the features or regions, and relationships among them, are computed. These may include properties based on gray level, color, or texture, as well as geometric properties (size, shape, etc.). Resegmentation may also take place at this stage in order to obtain features or segments having desired properties. The results of this stage can be represented by a graph-like relational structure in which the nodes correspond to features or regions, and the node or arc labels correspond to property or relation values.

c) "Models" for the objects to be recognized are also represented in relational structure form, where the labels contain information about the conditions that the property or relation values must satisfy in order for the object to be of a given type. Object recognition can then be based on determining whether the observed relational structures satisfy these conditions.

CH2141-0/85/0000/0002$01.00 © 1985 IEEE

INTERMEDIATE-LEVEL VISION

The block diagram of a generic intermediate-level vision system for three-dimensional scene analysis is more complicated, as we see from Figure 2. Here the main stages are as follows:

a) Our first task is to "recover" information about the visible surfaces in the scene. This information can be regarded as "2-1/2 dimensional", since it relates only to visible surfaces, not to the full 3D content of the scene. A variety of techniques have been developed for inferring the illumination, reflectivity, and surface orientation at each visible surface point based on the gray level at the corresponding image point. Some of these techniques work directly from the gray level variations in the image, while others begin by extracting features (edges, texture primitives, etc.) from the image. [In this context, the feature map is sometimes called the "primal sketch".] The result of this stage can be thought of as a stack of "intrinsic images", each representing a different surface property (illumination, reflectivity, orientation); it is sometimes called the "2-1/2 D sketch".

b-c) We can now apply feature extraction and segmentation techniques to the 2-1/2 D sketch in order to detect various types of features (e.g., shadow edges, which are discontinuities in illumination; occluding edges, which are discontinuities in range, etc.) or to segment it into surface patches of various types. The results can be represented as a symbolic image. Properties of and relationships among the features or patches can then be measured and represented in a relational structure.

d) "Models" for three-dimensional objects are normally represented as object-centered, viewpoint-independent relational structures. To recognize objects in the image, we must find observed relational structures that are consistent with how pieces of the objects might appear from some viewpoint. This requires a constraint analysis process in which we verify that structures derived from the image are consistent with the "projections" of object structures.

HIGH-LEVEL VISION

In both of our vision system block diagrams, processes are performed in a fixed order, beginning with the input image and proceeding toward higher levels of abstraction, until a level is reached where comparison with models becomes possible. A system restricted to this type of operation is limited in its flexibility. It has no provision for evaluating the results at each stage and modifying the processing so as to improve these results.

Greater flexibility, at the cost of slower processing speed, would be achieved by an "expert" vision system in which knowledge about the scene domain, and the rules governing how this knowledge is used in analyzing the image, are explicitly represented, rather than being implicit in the design of the analysis procedures.

We will not attempt to draw a block diagram of an expert vision system, but we will make a few general remarks about the design of such systems.

a) Building an expert vision system cannot simply be done by somehow combining a standard "expert system" software package with a set of image processing operations. The knowledge needed in a vision system is primarily spatial knowledge, concerning the geometric properties of and relationships among objects. Much of this knowledge is very hard to verbalize and will not be easily expressible in any standard language or notation.

b) Knowledge should be formulated at the level of the scene and the imaging process, rather than at the level of the image. This allows the knowledge to be independent of viewpoint, occlusions, etc. Predictions about the image can be derived from knowledge about the scene and the imaging process, but reasoning in the reverse direction is much harder.

c) Knowledge should be used at the earliest stages of the image analysis process; in particular, it should be used in choosing feature extraction and segmentation techniques. An expert vision system should not operate on symbolic data only; it should be able to reason about raw images too.

d) An expert vision system should have at its disposal many different processes for feature extraction, segmentation, property measurement, etc. It should know about the behavior of these processes when they operate on images of a given type and should be able to choose the processes that have the greatest expected payoff (in terms of expected informativeness vs. computational cost) at a given stage of the processing and in a given part of the scene.

e) An expert vision system will have many types of knowledge at its disposal, and will need some method of selecting the appropriate pieces of knowledge to use at a given stage of the processing and in a given part of the image. This "indexing problem" will arise even in specific scene domains; it is not unique to "general-purpose" vision systems.

f) An expert vision system will need to employ massive parallelism, not just in processing images, but in reasoning about the entities extracted from the images. The large number of available processing paths will lead to a potential combinatorial explosion of inferences, and it will be necessary to develop methods of keeping this explosion under control.

These remarks suggest that the design of expert vision systems will present many challenges. However, future generations of vision systems will need to become increasingly "expert"; thus it is important that these challenges be met and successfully overcome during the coming years.

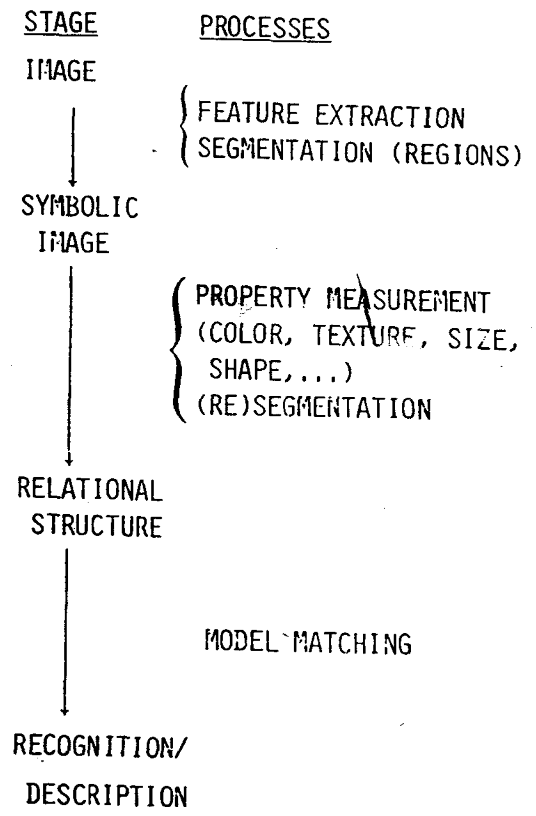

Figure 1: Block diagram of a low-level vision system.
(above)

Figure 2 (on right): Block diagram of an intermediate-level vision system.

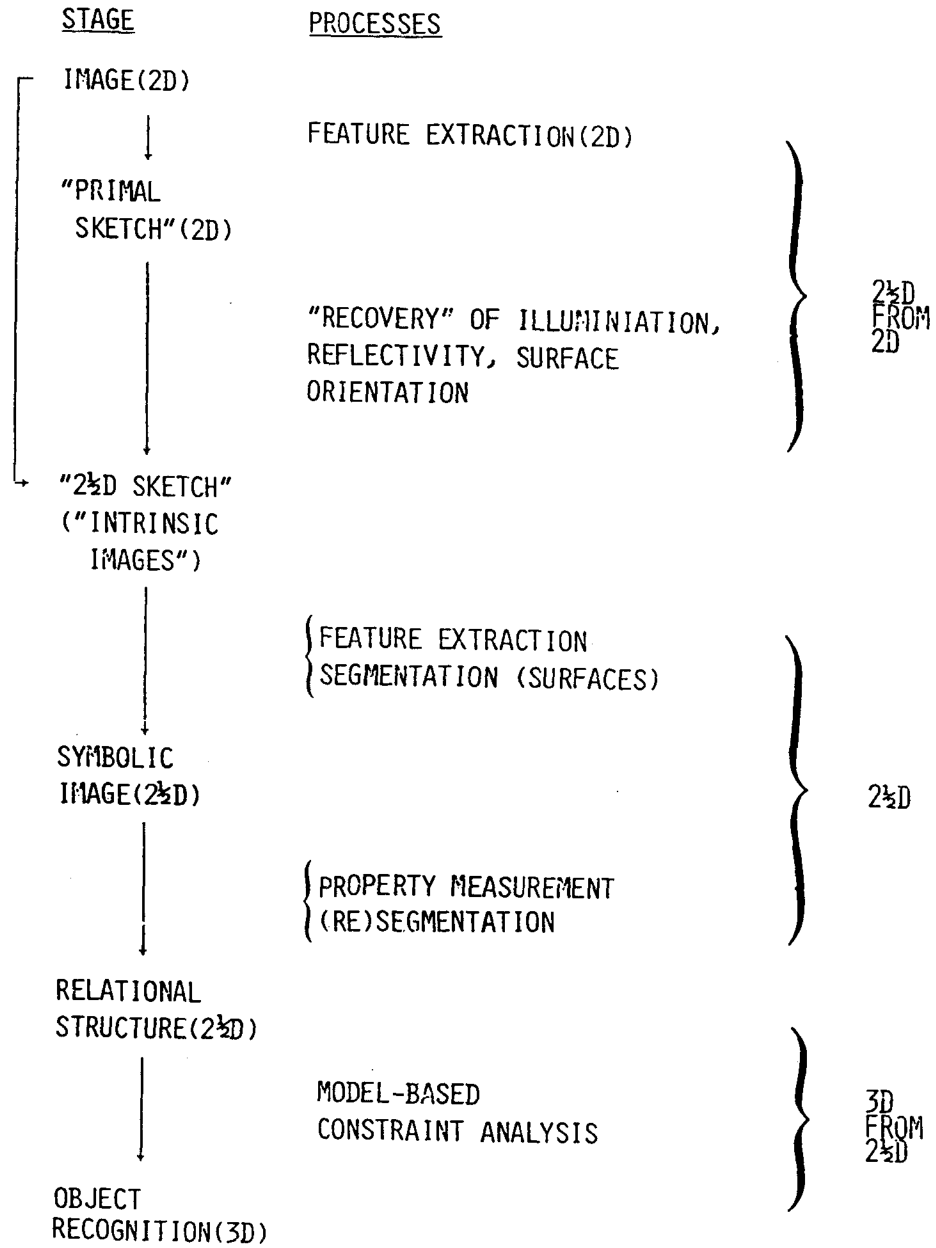

Cognitive Aspects of Reconstructability Analysis

Chairman
G. Klir
State University of New York-Binghamton

COGNITIVE ASPECTS OF RECONSTRUCTABILITY ANALYSIS

George J. Klir

Department of Systems Science
Watson School
SUNY Binghamton
Binghamton, New York 13901

ABSTRACT

The term "reconstructability analysis" (RA) has evolved as a generic term for all aspects of a methodology whose aim is to deal with the various problems associated with the relationship between systems perceived as wholes and their various subsystems (i.e., parts of the wholes). This paper is a conceptual overview of RA and its role in cognition. The focus is on characterizing problems that are studied under RA. Also outlined is a principle of inductive inference that is embedded in RA.

1. Introduction

An important feature of cognition--the process of knowing--is the ability to deal effectively with the relationship between parts and wholes. When perceiving a solid object, for example, we receive only its two-dimensional projection on the retina. Yet, we are able to reconstruct, using all relevant background knowledge, an image of the whole object that conforms to the real object to some degree. Similarly, when parts of English text are unreadable, we are often able to reconstruct the whole text or, at least, some of the unreadable parts. Hence, we are able to make ampliative inferences, that is inferences in which conclusions obtained contain more information than is contained in the premises.

At this time, we have little understanding of how the human mind makes ampliative inferences that involve the relationship between parts and wholes. We have, however, an organized collection of ideas, concepts, principles, and methods for formulating and dealing with the various questions associated with the whole-part relationship. This collection, which has been a subject of research since 1976[17], is now generally recognized under the name reconstructability analysis[22].

The most fundamental concept in reconstructability analysis (RA) is that of a system. In general, each system is conceptualized in RA as a set of variables that constrain each other in some specific way. It is understood that the variables of a system correspond to some real-world attributes, either as direct observables or as lagged observables, or are merely some abstract entities with no apparent real-world meaning (e.g. postulated internal variables or variables in purely mathematical systems). Although these distinctions between variables are significant (e.g. for characterizing dynamical aspects of systems), their affect on RA is only marginal.

Two complementary problems are involved in RA. In one of them, a set of subsystems (parts) of an unknown overall system (a whole) is given and the aim is to derive from the information in the subsystems as much information as possible regarding the overall system. This problem has been called the identification problem. The second problem is based on the assumption that an overall system is given; the aim is to determine which sets of its subsystems are adequate for reconstructing it, to an acceptable degree of approximation, solely from the information included in the subsystems. This problem has been referred to as the reconstruction problem.

The two problems of RA can thus be summarized as follows:

IDENTIFICATION PROBLEM

A set of subsystems (parts) $\longrightarrow$ Overall system (whole)

RECONSTRUCTION PROBLEM

An overall system (whole) $\longrightarrow$ "Best" sets of subsystems (parts)

The status of a system as either an overall system or a subsystem is not absolute. Any given system can assume the role of a subsystem (a part) in one context and the role of an overall system (a whole) in another context. This duality makes it possible to represent each overall system by a hierarchy of collections of subsystems, i.e., by a collection whose subsystems are also represented by collections of subsystems, whose subsystems are also ..., etc. Where we set the boundaries for what is to be considered the overall system and what are its subsystems is strictly a pragmatic question.

2. Identification Problem

The identification problem is based on the assumption that only information regarding subsystems of an overall system of concern is available. This partial information may not be sufficient to identify the overall system uniquely. Hence, one aim of the identification problem is to determine the class of all overall systems implied

CH2141-0/85/0000/0006$01.00 © 1985 IEEE

by the given information, to which the actual overall system must belong. This class is called a reconstruction family[3,10].

Given information regarding subsystems of an unknown overall system, the determination of the reconstruction family is a matter of deductive and, thus, nonproblematic inference. It amounts to formulating and solving appropriate algebraic equations that characterize, within the formal framework employed, the relation between the unknown overall system and the given subsystems. These equations are always constrained by special requirements of the formal apparatus employed. When probability theory is used, for example, the unknown probabilities of overall states must be nonnegative numbers that add to one.

Methods for determining the reconstruction family have been developed for systems conceptualized within the frameworks of probability theory[3,12] and possibility theory[10]. For probabilistic systems, a method for measuring the cognitive content of individual substates is now available; it was developed by Jones and is described in this volume[16].

A reconstruction family may contain one or more overall systems, or it may be empty. When it contains only one system, the identification is perfect. When, on the other hand, it contains more then one system, the identification remains undecided. The larger the reconstruction family, the less precise is the identification of the overall system. In general, when new information is added, the reconstruction family becomes smaller and, consequently, the identification becomes more precise.

When the reconstruction family is empty, the given information regarding subsystems is inconsistent. Two kinds of inconsistency among subsystems are distinguished in RA--local and global inconsistency. A pair of subsystems that share some variables is locally inconsistent if the shared variables are required to have different properties in the two subsystems. A collection of subsystems is globally inconsistent if it is locally consistent and, yet, the reconstruction family is empty.

Local inconsistencies may be (and usually are) caused by the fact that the subsystems themselves are not perfect models of reality and are developed from or validated by different data sets, each of which is limited, incomplete, imprecise, or otherwise imperfect in some specific sense. This means that local inconsistencies are a result of our ignorance regarding the subsystems involved and, consequently, it is meaningful to try to resolve them by a suitable method; this issue is discussed by Mariano in this volume[23].

A global inconsistency is more serious. It usually indicates that the collection of subsystems is ill-conceived; it is a mathematical artifact that has no meaning in real world. Although global inconsistencies can also be resolved in principle, it is not clear how to justify any particular formulation of the problem.

In addition to the determination of reconstruction families, the concept of an identifiability quotient has been introduced in RA, by which the ability to identify a unique overall system from information regarding its subsystems is measured[3,10]. In general, the identifiability quotient reflects the relative size of the reconstruction family; the smaller the size, the larger the value of the identifiability quotient.

All members of a reconstruction family are hypothetical overall systems implied by the given subsystem information, but it is not known which of them is the actual overall system. If we must select one of these hypothetical overall systems (e.g., as a basis for making a decision and taking an appropriate action), we need some justification for choosing one of them over the others. That is, we must resort to some justifiable principle of ampliative inference. One principle, well justifiable on epistemological grounds, can be derived for each formal framework from the requirement that the chosen overall system be unbiased, i.e., implied solely by the given information about subsystems. This unique member of the reconstruction family is called an unbiased reconstruction.

To determine the unbiased reconstruction requires that all but no more information than available be used. For probabilistic systems, this requirement leads uniquely to the well justified principle of maximum entropy[5,11 24]. Similar principle has recently been developed for possibilistic systems[4]. A number of computationally efficient algorithms are now available for making these principles operational[3,4,13-15].

The unbiased reconstruction is not the only justifiable choice from the reconstruction family. On pragmatic grounds, for example, we may require to choose such a member of the reconstruction family for which the maximum possible error is minimized; a member with this property is usually referred to as the least risk reconstruction or min-max reconstruction.

3. Reconstruction Problem

In the reconstruction problem, a system conceived as an overall system is given and the aim is to determine from among the many collections of its various subsystems those which allow us to reconstruct the overall system to an acceptable degree of approximation. The primary reason for trying to represent a given overall system by a set of its subsystems is to reduce its complexity and, consequently, to improve its manageability. For example, it is easier to monitor four sets of three physiological variables of a patient independently of each other than to monitor all variables (say 10 of them) simultaneously during a difficult surgery, when decisions about appropriate actions must be made quickly.

In dealing with the reconstruction problem, all meaningful sets of subsystems of a given overall system must be considered, evaluated, and compared in terms of their complexities and reconstruction capabilities. There are 2^n of subsystems and 2^{2^n} sets of subsystems of an overall system with n variables. Not all of these sets, however, must be considered in the reconstruction problem. Indeed, we may disregard such sets that contain at least one element which is a subsystem of another element. Such an element, which is a sub-subsystem of the overall system, does not contribute any information that is not included in the larger element and, hence, it is totally redundant in the context of the reconstruction problem.

Set of subsystems of a given overall system that do not contain redundant subsystems are viewed as meaningful reconstruction hypotheses. They are ordered by a relation of refinement. Given two reconstruction hypotheses, X and Y, X is viewed as a refinement of Y if and only if for each subsystem in X there is a larger or equal subsystem in Y; the term "larger" is used here strictly in the sense of the sets of variables associated with the subsystems. The refinement ordering is only partial. It is known that the set of all reconstruction hypotheses for any overall system together with the refinement ordering form a lattice[2].

Given a set of reconstruction hypotheses of an overall system of concern, we need to evaluate how much information about the overall system is contained in each of them. Since such evaluations are based on the hypothetical overall systems reconstructed from the various hypotheses, we must insure that the reconstruction method used utilizes all information available in each of these hypotheses. At the same time, however, we must be sure that no additional and unsupported (i.e. biasing) information is used in deriving the reconstructed system. Hence, we must determine the unbiased reconstruction for each relevant reconstruction hypothesis and, then, compare it with the actual overall system. The comparison is expressed in terms of a suitable distance function, preferably one based on the relevant measure of information[1,8,9].

Our aim in the reconstruction problem are to maximize refinement and, at the same time, minimize distance. However, these two aims conflict with each other. Any increase in the refinement implies that the distance increases or, at best, remains the same. These considerations leads to the following formulation of the reconstruction problem.

Given an overall system X and a set H_X of its reconstruction hypotheses, let $\overset{r}{\leq}$ and $\overset{d}{\leq}$ denote the refinement ordering and distance ordering on H_X, respectively, let $\overset{\alpha}{\leq}, \overset{\beta}{\leq}, \ldots$ denote some additional (optional) preference orderings on H_X, and let a joint preference ordering $\overset{*}{\leq}$ on H_X be defined by

$$g \overset{*}{\leq} h \Leftrightarrow g \overset{r}{\leq} h \text{ and } g \overset{d}{\leq} h \text{ and } g \overset{\alpha}{\leq} h \text{ and } g \overset{\beta}{\leq} h \ \ldots$$

for each pair $g,h \in H_X$. The solution set S_X of the reconstruction problem for X is a subset of H_X such that $S_X = \{h \in H_X \mid (\forall g \in H_X)(g \overset{*}{\leq} h \Rightarrow h \overset{*}{\leq} g)\}$, i.e., reconstruction hypotheses in S_X are either equivalent or incomparable in terms of the joint preference ordering $\overset{*}{\leq}$.

The reconstruction problem involves some difficult computational issues. They are associated primarily with the search through relevant reconstruction hypotheses, whose numbers grow extremely rapidly with the number of variables in the analyzed overall system. While there are only 18 reconstruction hypotheses for three variables and 166 for four variables, there are 7,579 of them for five variables, almost eight million for six variables, and about 2.4×10^{12} for seven variables. Various search strategies have been developed[2,21], but it is beyond the scope of this paper to cover these technical details.

4. Reconstruction Principle of Inductive Inference

Extensive simulation experiments have been performed since 1976 to determine some fundamental characteristics of RA[7,20]. Some of these experiments and results obtained are described by Hai in this volume[6]. Among other results, the experiments confirmed, at least partially, a novel principle of inductive inference, which I proposed in 1981[18]. Since it is embedded in RA, the principle is now referred to as the reconstruction principle of inductive inference.

The reconstruction principle of inductive inference proceeds in two stages. In the first stage, an overall constraint among the variables of concern is estimated from the available data by using the usual principles of inductive reasoning. For example, probabilities of overall states of the variables are estimated by the principle of maximum entropy. The second stage consists of three steps:

(i) superior reconstruction hypotheses (those with minimum information distance) are determined for the overall system at the various levels of the refinement lattice;

(ii) beliefs of various degrees that these superior hypotheses reflect the actual reconstruction properties of the variables involved are formed on the basis of relevant experimental characteristics of RA;

(iii) the given overall constraint is supplemented with (or replaced by) the constraints reconstructed by the superior reconstruction hypotheses with the respective degrees of belief at the individual refinement levels.

While using only the information included in the available data, this two-stage method allows us to include in the estimated overall constraint certain features (e.g. some overall states) which are not directly derivable from the data. Hence, it allows us, for instance, to predict or retradict, with a specific degree of belief (credibility), certain states of the variables that are not included in the data available at the time of making the prediction or retrodiction. It is the capability of RA to determine the superior reconstruction hypotheses that makes it possible to produce this novel information. While this information is not available (by definition) in the data explicitly, it is implicitly incorporated (encoded) in their reconstruction properties. RA enables us to determine these reconstruction properties and, consequently, it enables us to recover (decode) this implicit information.

It might be interesting and potentially fruitful to investigate one of the least understood functions of the human mind--intuition--in terms of the reconstruction principle of inductive inference.

5. Conclusions

Since the initial inception of RA in 1976[17], many conceptual mathematical and computational results have been obtained, which have made it possible to develop an appropriate computer software to deal with both the identification and reconstruction problems. The most comprehensive and up-to-date summary of RA--its aims, results, and open problems--can be found in a recent paper I coauthored with Eileen Way[22].

Within the broader context of systems problem solving, RA represents one important category of methodological tools--those for dealing with the interplay between wholes and parts. It is presented in this context in my recent book[19].

In general, RA provides us with methodological tools for reasoning and decision-making in the face of evidence that is fragmentary, incomplete, imprecise, locally inconsistent, or otherwise imperfect (identification problem). It also provides us with tools for making systems more manageable by breaking them into appropriate sets of subsystems (reconstruction problem). There is no doubt that some tools of this sort are essential for cognition and, as such, are embedded in one form or another in the human mind. Further advances in RA may thus help us not only to improve our capabilities of dealing with various practical problems, but also to advance ones insights into the operation of our minds.

REFERENCES

[1] Aczel, J. and Z. Daroczy, On Measures of Information and Their Characterizations. Academic Press, New York, 1975

[2] Cavallo, R.E. and G.J. Klir, "Reconstructability Analysis of multi-dimensional relations: a theoretical basis for computer-aided determination of acceptable systems models." Int. J. of General Systems, 5, No.3, 1979, pp. 143-171.

[3] Cavallo, R.E. and G.J. Klir, "Reconstructability analysis: evaluation of reconstruction hypotheses." Int. J. of General Systems, 7, No. 1, 1981, pp. 7-32.

[4] Cavallo, R.E. and G.J. Klir, "Reconstruction of possibilistic behavior systems." Fuzzy Sets and Systems, 8, No. 2, 1982, pp. 175-197.

[5] Christensen, R., Entropy Minimax Sourcebook. Vol I: General Description. Entropy, Lincoln, Mass., 1981.

[6] Hai, A., "On the development of a knowledge base for General Systems Problem Solver." This volume

[7] Hai, A. and G.J. Klir, "An empirical investigation of reconstructability analysis." Int. J. of Man-Machine Studies, 22, 1985 (in press).

[8] Higashi, M. and G.J. Klir, "Measures of uncertainty and information based on possibility distributions." Int. J. of General Systems 9, No. 1, 1983, pp. 43-58.

[9] Higashi, M. and G.J. Klir, "On the notion of distance representing information closeness: possibility and probability distributions." Int. J. of General Systems, 9, No. 2, 1983, pp. 103-115.

[10] Higashi, M., G.J. Klir, and M.A. Pittarelli, "Reconstruction families of possibilistic structure systems." Fuzzy Sets and Systems, 12, No. 1, 1984, pp. 37-60.

[11] Jaynes, E.T., "Where do we stand on maximum entropy?" In: The Maximum Entropy Formalism, edited by R.L. Levine and M. Tribus, MIT Press, Cambridge, Mass., 1979, pp. 15-118.

[12] Jones, B., "Determination of reconstruction families." Int. J. of General Systems, 8, No. 4, 1982, pp. 225-228.

[13] Jones, B., "Determination of unbiased reconstructions." Int. J. of General Systems, 10, No. 2, 1984.

[14] Jones, B., "A greedy algorithm for a generalization of the reconstruction problem." Int. J. of General Systems, 11, No. 1, 1985 (in press).

[15] Jones, B., "Reconstructability analysis for general functions." Int. J. of General Systems, 11, No. 2, 1985 (in press).

[16] Jones, B., "The cognitive content of system substates." This volume.

[17] Klir, G.J., "Identification of generative structures in empirical data." Int. J. of General Systems, 3, No. 2, 1976, pp. 89-104.

[18] Klir, G.J., "On systems methodology and inductive reasoning: the issue of parts and wholes." General Systems Yearbook, Vol. 26, 1981, pp. 29-38.

[19] Klir, G.J., Architecture of Systems Problem Solving. Plenum Press, New York, 1985.

[20] Klir, G.J. and H.J.J. Uyttenhove, "On the problem of computer-aided structure identification: some experimental observations and resulting guidelines." Int. J. of Man-Machine Studies, 9, No. 5, 1977, pp. 593-628.

[21] Klir, G.J. and H.J.J. Uyttenhove, "Procedures for generating reconstruction hypotheses in the reconstructability analysis." Int. J. of General Systems, 5, No. 4, 1979, pp. 231-246.