

JOHN P. GALLAGHER

KNOWLEDGE SYSTEMS FOR BUSINESS

INTEGRATING EXPERT SYSTEMS & MIS

Knowledge Systems --- for Business

Integrating Expert Systems
& MIS

John P. Gallagher



PRENTICE HALL, *Englewood Cliffs, New Jersey 07632*

Library of Congress Cataloging-in-Publication Data

GALLAGHER, JOHN P. (date)

Knowledge systems for business: integrating expert systems and MIS / John P. Gallagher.

p. cm.

Includes index.

ISBN 0-13-516551-2

1. Management information systems. 2. Expert systems (Computer science) I. Title.

T58 .6.G336 1988

658.4'038--dc 19

The publisher offers discounts on this book when ordered in bulk quantities.

For more information, write:

Special Sales/College Marketing
Prentice-Hall, Inc.
College Technical and Reference Division
Englewood Cliffs, NJ 07632



© 1988 by Prentice-Hall, Inc.

A Division of Simon & Schuster

Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-516551-2

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

To Robin

•

•

Preface

When the impact of knowledge systems on business practice is reduced to its simplest form, a relatively straightforward and simple phrase remains: *What is feasible has changed.*

There are many ways to marginally influence the efficiency, effectiveness, and profitability of any business unit. However, there are few opportunities to fundamentally alter the competitive structure or competitive scope within entire industries. In the past few years, information technology has increasingly brought about these types of changes. This history of impact, coupled with an impressive portfolio of success stories, has heightened management's interest in the strategic, competitive applications of information technology. Whenever what is feasible changes, significant business opportunities and problems follow. The theme of this book is the integration of knowledge systems, not only with existing information systems, but with the existing organization that influences, designs, supports, and applies them.

The origins of this work are easily traced. For the past five years, I have served as Director of Computing at the Fuqua School of Business, Duke University. In this capacity, my efforts have much in common with corporate organizations labeled variously as "Emerging Technologies" or "Strategic Information Systems" groups. That is, much of my time has been spent exploring new computing technologies and assessing their relevance for research, for the graduate curriculum in management education, and for executive education programs.

If there is one lesson I have learned through this experience, it is that in order for any computing technology to affect the practice of management, a constellation of factors must exist to create a receptive context. As an agent of change, technology is a catalyst. Technology stimulates change because of its interaction with multiple factors in a complex environment.

This book is an effort to define some of the characteristics of the technology of knowledge systems in terms of the environment of management computing. It attempts to create a broader perspective that incorporates the motivations of individual business units, the nature of their objectives, management's orientation toward computing, and the "installed base" of mainstream computing technology and professionals.

My involvement with artificial intelligence began with my doctoral studies at the University of California, Santa Barbara, in 1974. I was employed on a research project in the Computer Systems Laboratory to develop artificial intelligence applications for teaching mathematics problem solving skills with computers. That experience resulted in a strong vision of both the potential of this technology to support human intellectual activity, and the complexity of the problems inherent in its development.

In the past fifteen years, the technology has matured and come to be better understood by individuals outside of the closely connected world of academic research. Just a short time ago, the first knowledge systems development tools that were designed to be used by individuals unfamiliar with, and largely uninterested in, the technology's academic origins were commercially offered. Regardless of whatever significance might be attached to such an offering, it raised an important question: *Could the developer/user of a knowledge systems development tool be sufficiently isolated from the complexities of the technology to be productive in the creation of applications that exploit its unique capabilities?*

For the past three years I have been teaching executive audiences and graduate business students with these tools. I am convinced that the question raised above has been clearly answered. The next round of questions relate to the more complex issues that will determine the degree to which these technologies can be effective in their impact on management practice. This book is an attempt to articulate some of these issues.

JOHN P. GALLAGHER

Contents

Preface	vii
1 Introduction	1
Using Information Technology to Gain a Strategic Competitive Advantage	6
2 Knowledge Systems Development Tools	15
Evolution from Techniques to Tools	17
Choosing Techniques and Tools for Knowledge Systems	19
The Nature and Impact of Available Tools	22
Separation of Knowledge and Reasoning	25
The Inference Process: Backward Chaining	29
Dealing with Uncertainty	33
The Inference Process: Forward Chaining	35
Frame-Based and Object-Oriented Systems	37
3 Evaluating Knowledge Systems Development Tools As Mainstream Business Computing Systems	41
Knowledge Systems and Mainstream Computing Hardware	44
Access to Data and Other Applications	47
Representing and Processing Information from Other Sources	50
The Support Environment for Applications Development	53
4 Some Sample Applications	57
A Load Control Advisor for the Small Electrical Utility Company	61
An Expert Data Analyst/Advisor for Consumer-Packaged-Goods Marketing	66
A Financial Statement Analysis Advisor	70

5 Knowledge Engineering: The Development of Intelligent Systems Applications	75
Traditional Structured Systems Development	77
Knowledge Engineering: Designing and Developing Expert Systems	78
The Holistic Approach to Systems Development	81
Representing the Structure of Knowledge	85
Automated Knowledge Acquisition: Rule-Induction Systems	89
6 A Review of Development Tools: M.1 from Teknowledge, Inc.	95
The Promotion Advisor Knowledge Bases	97
M.1 from Teknowledge, Inc.: An Overview	100
The Development Environment	101
Inferencing and Dealing with Uncertainty	104
Integration and Noninferential Processing	107
The Promotion Advisor in M.1	109
Summary	118
7 A Review of Development Tools: GURU from Micro Data Base Systems	119
Overview	121
The Development Environment	122
Inferencing and Dealing with Uncertainty	128
Integration and Noninferential Processing	130
The Promotion Advisor in GURU	133
8 A Review of Development Tools: The Expert System Environment from IBM	137
Overview	139
The Development Environment	140
Inferencing and Dealing with Uncertainty	144
Integration and Noninferential Processing	146
The Promotion Advisor in ESE	147
9 Forces Affecting the Future of Knowledge Systems in Business	153
The Need for Organizational Accommodation/Assimilation	157
The Need to Develop Supportive Technology	158
Uncertainty of the Application's Suitability for Knowledge Systems	159
Competitive Forces, the Acceptance of Risk, and the Future of Knowledge Systems	160
Appendix: Knowledge Bases for the Promotion Advisor Written in M.1, GURU, and Expert System Environment	163
Index	229

1

Introduction

Increasingly, popular management literature has focused on the application of artificial intelligence technology to the solution of management problems. Most discussions in this area have addressed either the technology of various programming techniques or success stories of large-scale corporate profitmaking through their application.

Several interacting factors contribute to corporate interest in this technology. First, what were once remote, artificial intelligence programming techniques have been significantly redesigned and presented to the business computing market in the form of applications development tools. Recent advances in computing hardware have provided a suitable foundation for the delivery of these tools on new generations of “mainstream” business computers. These hardware and software technologies promise the kinds of power and ease-of-use that have been so successful in other applications development environments like spreadsheet modeling and data base management systems. The potential of such tools to open broad new territories to computing applications is very exciting and has captured management’s attention.

Second, there is a developing perception of both information and knowledge as corporate assets. New computing technologies promise to make knowledge that has been implicit in the behavior of decision makers explicit in a machine-usable form. By codifying knowledge, managers make it a manageable asset, continuously available to their organization. The potential integration of machine-usable knowledge with machine-readable information promises to carry the information age into the era of “knowledge management.”

Finally, the role of information technology is being redefined as a competitive weapon. This is in sharp contrast to internal applications of technology to gain increased efficiencies. As business planning looks to technology to help differentiate product offerings, new technologies like expert systems take on added significance. An interesting series of transitions has taken place in the perception of the role of computing as it has moved from data processing to information systems to management information systems and to strategic information systems.

The cumulative effect of these forces has created an atmosphere responsive to the promise of knowledge-oriented computing systems. The promise appears to be quite real, and the challenge, of course, is to exploit it to advantage. Just how this is to be accomplished depends on a number of factors including individual

lines of business, competitive environments, existing attitudes toward information systems management, orientations toward risk, and so on, in addition to the technology itself.

Chapter 2 includes a more detailed discussion of the distinctions among artificial intelligence, expert systems, and knowledge systems. In general, however, artificial intelligence is the superordinate concept in this set. As a field, artificial intelligence is concerned with computing technologies that allow machines to accomplish tasks previously believed to be the exclusive domain of human intelligence. There are many smaller fields within artificial intelligence, such as natural language understanding, vision processing, speech generation, robotics, and the like. One of these areas, included in the larger field of artificial intelligence, is expert systems.

Expert systems, as a field, emerged fairly early in the development of artificial intelligence. It resulted from efforts to model human problem solving with computing technology. As a research paradigm, it became useful to model individual, exemplary human problem solvers within highly specific application areas. The modeling of these individuals resulted in the term "expert system." Since these early efforts, generalized techniques have evolved for representing human knowledge in a machine-usable form and applying it to solve a defined class of problems. Even though the techniques evolved from the early "expert system" research efforts, they are now being applied to modeling and using knowledge from sources other than individual "experts." Many applications are being based on composite sources of knowledge, such as more than one individual, books, manuals, regulations, reports, published procedures, and so forth. These systems employ the same underlying programming technologies as expert systems, but are based on a broader source of knowledge. As such, they are referred to by the more general term "knowledge systems." Chronologically, the emergence of knowledge systems followed expert systems. However, in terms of superordinate and subordinate relations, knowledge systems are a subordinate concept to artificial intelligence, and a superordinate concept to expert systems.

The argument made here is that although it does make sense under specific circumstances for management to initiate in-house artificial intelligence research and development efforts, the vast majority of this technology's impact will come from the growing number of tools designed for knowledge systems application development by end-users and information systems professionals. At present, the market has experienced the first wave of such tools. Although they range in flexibility and power, they share some common traits that are significant in determining their importance as a class. First of all, they run on the hardware systems familiar to the world of corporate computing. Second, they are designed to conform to the pattern of other, popular development tools: they require a minimum of ancillary study and knowledge for their effective application and use by end-users and/or systems professionals. Finally, whereas they represent only a subset of the techniques that define the broader field of artificial intelligence, these techniques have been proven to be robust and appropriate for a broad range of applications that typify managerial decision making.

The challenge is to integrate these tools with the mainstream of business computing to achieve management's objectives. With this perspective, this book is about practical matters in the selection, development, management, and maintenance of knowledge-based systems with this new generation of development tools. A minimum of space is allotted to the more esoteric issues from the field of artificial intelligence. Discussed in their place will be the more germane issues of how knowledge systems technology fits within the context of corporate computing, how the nature of managerial computing is changing, and how to plan for, prototype, implement, and manage new generations of knowledge systems for management. In short, this book is about "getting on with it."

There is, and will continue to be, a considerable variance in available knowledge systems application development tools. Vendors have provided a wide range of functions and features in their attempts to distinguish their products. However, the underlying technologies are relatively straightforward. The common technological base for these products will be discussed in a generic fashion. Later, however, specific tools will be discussed in some detail to better describe the value and applicability of some of the more important distinguishing characteristics of present offerings.

In most instances, organizations with an interest in applying knowledge systems find themselves faced with a number of questions, only some of which are directly related to the technology itself. Examples include

Who will provide project leadership?

That is, to what extent is the development of knowledge systems fundamentally different from existing applications? Under what circumstances will current information systems professionals be able to provide support and leadership in systems development? Is the very nature of the technology sufficiently distinct to require new approaches to development? Or are the applications themselves different enough to thwart existing procedures for systems development? Who is best able to identify and evaluate potential application areas?

Who will implement and maintain systems as they are developed?

Regardless of the source of leadership in project identification and orientation, there are questions as to whether or not knowledge-oriented computer systems are sufficiently similar to other computer applications to be written and maintained by the same professionals who coded the payroll systems. That is, is it necessary to compete for and hire highly trained graduates of artificial intelligence departments to investigate these technologies and develop and maintain applications? Or can existing organizational units learn to apply these techniques and develop and maintain knowledge systems? Will knowledge systems follow the pattern of other end-user, workstation-oriented development tools because of the close ties of potential applications to the knowledge of functional area workers? Or might we also see efforts to centralize and explicitly manage corporate knowledge systems in a fashion similar to centralized data processing and information systems development?

What hardware and software systems should be purchased?

Most of the business applications of knowledge systems that are widely discussed have been developed with hardware and software systems specifically created to support artificial intelligence research. Is it important to invest the time and capital to introduce a new

computing technology? If so, how substantial a commitment is required? Can these highly specialized technologies be integrated with existing information systems? Or is there a danger in supporting parallel technologies? One of the reasons for the increasing interest in knowledge systems is the emergence of applications development tools. Their promise is to reduce the development cycle and operate in the computing mainstream alongside familiar hardware systems. Are these empty promises, or has a significant transition in systems development technology taken place?

Another characteristic of the approach taken here is a strong belief that these questions should be addressed in the larger context of corporate computing as it currently exists, and as it appears to be evolving. Each of the general questions raised above, and the multitude of smaller, related questions that emerge like layers of an onion, cannot be answered by understanding only the nature of these new tools. The apparently alien nature of the technology should not detract from the fact that its management is a management issue that must be understood in the larger context of the corporate computing mix.

USING INFORMATION TECHNOLOGY TO GAIN A STRATEGIC COMPETITIVE ADVANTAGE

Many of the issues related to knowledge systems development and management are closely tied to the changing nature of corporate computing itself. Recently, a high level of attention has been paid to the larger issues of strategic corporate computing and changes in the nature and role of information management.

Recent discussions of the nature of corporate computing have resulted in a variety of related conceptual models of the use of information technology as a strategic competitive force. Although several of these exist, one of the more general and useful, described by James Cash, is the "three era model." This model creates an image of overlapping categories, or types of corporate computing. These types of computing applications are distinct, have appeared sequentially, and have, therefore, been designated as "eras." The primary characteristics of each era, including its administrative framework, primary clients, and sources for justifying systems development, are depicted in Figures 1-1 through 1-3.

Era I, beginning in the early 1960s, focused on the dominant, transaction-oriented, backbone corporate information systems like those for order entry, personnel records, payroll, and sales data. Because of the time frame for their development, they were generally written in-house by emerging data processing departments in COBOL or some other third-generation language.

Large data processing departments have developed over the past 25 years to oversee the creation and maintenance of such systems. The resulting administrative framework has been characterized as a "regulated monopoly," in that the data processing department, in concert with other organizational units, has determined priorities for which systems are developed, modified, or placed in the ever-lengthening queues. The primary clients for Era I systems are either the corporate organization as a whole, or large administrative units. Justification for prioritizing Era I

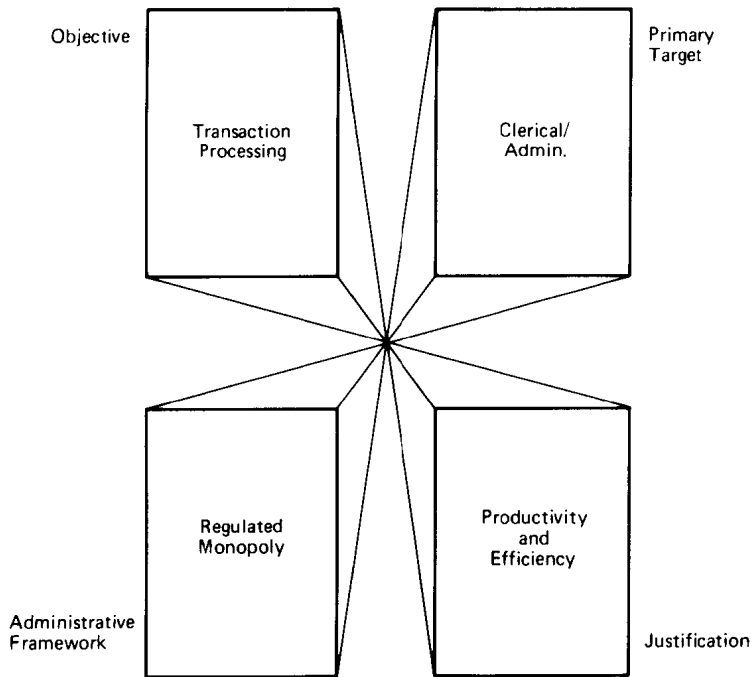


Figure 1-1 Era I




Figure 1-2 Era II




Figure 1-3 Era III

systems development has been based on the productivity-efficiency, cost-benefit model. Just how the management of these centralized information systems has evolved in many corporations is a critical factor in the enterprise's response to the challenges of Eras II and III.

With the advent of the early, stand-alone personal computer, beginning about 1980, Era II corporate computing was introduced. Whereas Era I computing focused on the processing of the organization's records, Era II has as its client individual managers and their needs for decision support. Numerous accounts are available of sophisticated information management and decision support systems developed either by, or under the close supervision of, the end-user. In many cases, these end-user systems have been comparable or superior in function to systems for which corporate MIS/DP departments had projected multiyear development cycles. For most of these applications the result is not a system of the type that the professional would develop and support, but individual managers have been willing to sacrifice these features for quick prototyping, rapid development, and the sense of self-control characteristic of personal computing solutions.

One result of Era II computing has been the development of small-scale Era I-type systems that have been waiting with a low priority in the MIS queue. More importantly, however, entirely new forms of computing, designed specifically for managerial decision support, have emerged. These include the variety of spreadsheets and other modeling tools that had no real counterpart in the mainframe MIS

environment. Era II applications have had a powerful impact on managerial views of computing applications, creating a climate of high expectation for the competitive applications of computing in general, and for new computing technologies in particular.

If one were to take a marketing perspective, what occurred in the transition from Era I to Era II has a lot to do with the concept of “sampling” or the use of “trial size containers.” Management, for relatively low cost, has now sampled a new style of support. A very competitive industry has developed to supply this lucrative market with even more effective and promising technology. The offerings of this industry have created a new market that spends significant amounts of time anticipating upcoming products. In a sense, the success of the current generation of personal computing products has created high expectations for future products and a predisposition to believe that new technology will continue to enhance managerial effectiveness. This attitude on the part of the managerial computing market has a great deal to do with the enthusiasm currently expressed toward knowledge systems technology.

Era II applications, for the most part, have not been under the control of a central information systems organization. It seems fair to characterize the administrative framework for Era I transaction-oriented systems as monopolistic, whereas Era II applications have been more akin to free market structures. The means of justifying Era II expenditures has largely been defined as managerial effectiveness as opposed to the administrative efficiency arguments used to justify Era I systems. This shift in the basis for the justification of systems development has been a contributing factor to the interest in knowledge systems technology as well.

Management has experienced a great deal of benefit from its introduction to Era II computing, owing to its unique attention to managerial decision support applications. This has resulted in a newly perceived relevance for the application of computing to general management. The apparent utility of knowledge systems applications for managerial decision making place them in this same context, as tools intended to extend and leverage managerial activity.

The third era of computing is closely tied to the concept of strategic competitive analysis. In its briefest form, this application of computing is concerned not with the general objective of making the internal, organizational operations of a business unit more efficient, nor with making the individual managers and/or management teams more effective. Rather, the focus is on using information technology to gain a competitive advantage.

Figure 1-4 describes some of the fundamental concepts of the competitive forces model as presented by Michael Porter. The competitive forces model is a popular and appropriate conceptual framework for clarifying the unique characteristics of Era III computing. In this model, such forces as potential new entrants, the bargaining power of buyers, the threat of substitute products/services, and the bargaining power of suppliers join with traditional intra-industry rivalry to act on any strategic business unit. Figure 1-5 adds to the model several approaches to reducing the impact of these competitive forces. For example, to reduce the force of potential new entrants to the market, the business unit in question can produce