
INTELLIGENT DATABASES

**Object-Oriented, Deductive
Hypermedia Technologies**

Kamran Parsaye

Mark Chignell

Setrag Khoshafian

Harry Wong

INTELLIGENT DATABASES

**Object-Oriented, Deductive
Hypermedia Technologies**

Kamran Parsaye

Mark Chignell

Setrag Khoshafian

Harry Wong



WILEY

John Wiley & Sons, Inc.

New York • Chichester • Brisbane • Toronto • Singapore

Copyright © 1989 by John Wiley & Sons, Inc.

All rights reserved. Published simultaneously in Canada.

Reproduction or translation of any part of this work beyond that permitted by Section 107 or 108 of the 1976 United States Copyright Act without the permission of the copyright owner is unlawful. Requests for permission or further information should be addressed to the Permissions Department, John Wiley & Sons, Inc.

Library of Congress Cataloging-in-Publication Data:

Intelligent databases.

Bibliography: p.

1. Data base management. 2. Expert systems (Computer science) I. Parsaye, Kamran.

QA76.9.D3I548 1989 005.74 88-33948

ISBN 0-471-50346-0

ISBN 0-471-50345-2 (pbk.)



PREFACE

Anthropologists have argued over whether humans became smart because they walked upright, because they had to kill other animals without having the powerful jaws and claws needed for the job, or because they learned how to communicate using a complex form of articulation and vocalization. However, lurking behind our definitions of humans as tool users and language users is the prior concept of humans as information users. Information and its use allow us to build and use tools and to express ourselves in language.

It is often remarked that we are living in the information age, yet we are struck by the fact that many existing technologies (1) are under-utilized and (2) tend to be treated in isolation. In the past few years, major advances have been made in several fields of critical relevance to information technology. These advances have been significant, but fragmented. Bringing these separate technical leaps into a unified framework will benefit people involved in all of the contributing disciplines by facilitating cross-disciplinary activity and by pointing out new avenues for both research and application.

This book is about the emerging intelligent database technology that will have a dramatic impact on the way we think and work. Intelligent databases will greatly expand our capabilities as information users. The phenomenon of the shrinking world is reaching its natural conclusion (i.e., McLuhan's global village), and intelligent databases are the means by which we may store and retrieve the information and expertise of this global village.

We may have to discard some familiar intellectual habits and presuppositions along the way. For instance, we are accustomed to thinking about information as passive,

residing in books and documents. Similarly, we often use libraries because we want to get at the information contained in their books, and we use databases as retrieval systems to find "answers." These "answers" are not always predefined in the system and may have to be created "on the fly," depending on the particular characteristics of the user and his or her current interests. Intelligent databases will change how we do research, how we look for ideas, how we make decisions, and how knowledge is transmitted.

In writing this book, we began with the premise that the database is at the heart of any information system and then sought to combine databases with a variety of tools that would make them more intelligent. In so doing, we have developed the model for intelligent databases that is described in this book.

The intelligent database model is based on five information technologies:

- Databases
- Object-Oriented Programming
- Expert Systems
- Hypermedia
- Text Management

Each of the four authors of this book has expertise in one or more of these areas. The book represents a balance between the need to provide sufficient background in each of the five technologies while also including new material on intelligent databases.

Databases and expert systems have been extensively discussed elsewhere, but the material in Chapters 2 and 4 distills the key concepts, thus allowing this book to be a self-contained treatment of intelligent databases. Although many texts on databases and expert systems are available, none relates them to the other technologies discussed here.

We also felt the need to provide clear expositions of object-oriented programming, hypertext, and text management. In addition to covering the major points made elsewhere in a number of scattered texts and journals, our treatment of these topics extends them where necessary so as to provide the tools for constructing intelligent databases.

The eight chapters of the book provide a complete and self-contained exposition of intelligent databases. Chapter 1 begins by defining intelligent databases and places them in the context of evolving information technology. The next five chapters present the key concepts for each of the five contributing technologies. In its discussion of databases, Chapter 2 focuses on relational and semantic data models in particular. Chapter 3 discusses the major issues of object orientation, including encapsulation,

inheritance, and object identity; it also reviews existing object-oriented languages and databases. Chapter 4 introduces the expert-system building blocks of knowledge representation and inference, and describes the process of building expert systems. The chapter concludes with a discussion of how expert systems and databases may be combined to form deductive retrieval systems.

Chapter 5 provides a thorough introduction to hypertext, relating it to the knowledge representation methods described in Chapter 4 and showing how hypertext networks may be linked to conventional databases. Chapter 6 reviews the extensive topic of text management, focusing on these elements most relevant to the construction of intelligent databases. Chapter 7 elaborates on the themes developed in Chapters 1 through 6. A formal model for intelligent databases is described that incorporates the five technologies described earlier. The main points of the earlier chapters are summarized in Chapter 8.

While writing this book, we constructed a prototype of the Fortune Finder intelligent database described in Chapter 7, thus demonstrating the viability of our intelligent database model. The material presented in this book represents the fruits of research and development that have been carried out by universities and industry over the past few years. In particular, we would like to acknowledge the contribution of Project Jefferson at the University of Southern California in demonstrating how hypermedia can be used for information management.

We write the book with several types of readers in mind. Those who want to master the topic of intelligent databases in its entirety should read all the chapters of the book in the order they are written. Database specialists who want to update their technical knowledge should skim Chapter 2 before focusing on Chapters 3, 4, 5, and 7. Administrators and managers should begin by reading Chapters 1 and 6, followed by Chapter 7. They should then return to the more technical material provided in such chapters as 4 and 5. Readers who have a good grasp on topics such as object orientation and expert systems, but who want to know what intelligent databases are, should concentrate, at least initially, on Chapters 1 and 7. Conversely, the book can also be used to review individual topics such as expert systems or hypertext. We see it being used in this way, for instance, in graduate-level college seminars on information technology or expert system applications.

All of the techniques described in this book can be implemented on a microcomputer such as a Macintosh II or a 286 or 386 class machine, and most of the techniques have already been implemented. In many ways, this book describes what is needed to build a shell for intelligent databases, much like the shells that are used to construct expert systems. In the same way that knowledge engineering of expert systems can be a difficult task, even when one has the expert system shell, so our experience has shown us that intelligent database engineering is a difficult task. However, we have no doubt that the benefits of intelligent databases will lead to them becoming a firmly established information technology.

We would like to thank those people who have helped us in preparing this book: Jenny Ghielmetti, Sandra Chignell, Silva Khoshafian, Linda Wong, Bob Blum, Bob Fraser, Joe Goguen, Lee Jaffe, Rei-Chi Lee, Diana Lin, Phil Smith, David Thompson and Gio Wiederhold. In preparing this book, we also benefited from discussion and interactions with a wide variety of colleagues, including the members of the Project Jefferson Team at USC.

KAMRAN PARSAYE
MARK CHIGNELL
SETRAG KHOSHAFIAN
HARRY WONG

*Los Angeles and Walnut Creek, California
January, 1989*

CONTENTS

1.

INTRODUCTION

- 1.1 Introduction / 1
- 1.2 Information and Information Glut / 6
 - 1.2.1 Precursors of the Information Society / 8
 - 1.2.2 The Hard-Copy Information Revolution / 8
 - 1.2.3 The Evolution of Information Technology / 11
 - 1.2.4 The Evolution of Early Computers and Databases / 14
- 1.3 Current Trends in Information Technology / 18
 - 1.3.1 Large-Scale Computing and Mass Storage / 19
 - 1.3.2 Computer Networks / 21
 - 1.3.3 Online Databases / 22
- 1.4 The Evolution of Intelligent Databases / 25
- 1.5 What Is an Intelligent Database? / 26
 - 1.5.1 Intelligence at the High-Level Tools Level / 27
 - 1.5.2 Intelligence at the User-Interface Level / 30
 - 1.5.3 Intelligence at the Database Engine Level / 31
- 1.6 Summary / 32

2.

DATABASE SYSTEMS

- 2.1 Introduction / 35
- 2.2 History and Evolution of Databases / 36

2.3	Modeling Data	/ 38
2.3.1	The Hierarchical Model	/ 39
2.3.2	The Network Model	/ 41
2.3.3	The Relational Model	/ 43
2.3.4	Entity-Relationship Data Model	/ 44
2.3.5	Semantic Data Models	/ 47
2.4	Levels of Database Architecture	/ 49
2.4.1	The Conceptual Level	/ 49
2.4.2	The Logical Level	/ 50
2.4.2.1	Schemas and Data Dictionaries	/ 50
2.4.2.2	Hierarchical and Network Schemas	/ 51
2.4.2.3	Relational Schemas	/ 52
2.4.3	The Physical Level	/ 53
2.4.3.1	Sequential Files	/ 54
2.4.3.2	Indexed Files	/ 54
2.4.3.3	Directly Accessed Files	/ 55
2.4.3.4	Hybrid File Organizations	/ 57
2.5	Relational Data Manipulation Languages	/ 57
2.5.1	The Relational Algebra	/ 57
2.5.1.1	Selection	/ 58
2.5.1.2	Projection	/ 59
2.5.1.3	Product	/ 59
2.5.1.4	Union and Difference	/ 60
2.5.1.5	Joining Tables	/ 61
2.5.2	SQL	/ 64
2.5.2.1	Simple SQL Statements	/ 65
2.5.2.2	Multitable SQL Queries and Subqueries	/ 67
2.5.2.3	Using SQL in a Host Language	/ 68
2.5.3	QBE	/ 69
2.6	Relational Database Design	/ 72
2.6.1	Capturing Requirements for Database Design	/ 72
2.6.1.1	Forms Approach	/ 72
2.6.1.2	Data Specifications	/ 74
2.6.1.3	Process Specifications	/ 75
2.6.2	Generation of Relational Structures and Normalization	/ 77
2.6.2.1	First Normal Form	/ 78
2.6.2.2	Second Normal Form	/ 78
2.6.2.3	Third Normal Form	/ 79
2.6.3	Process Transformation	/ 80
2.6.3.1	Flat-Form Processes	/ 80
2.6.3.2	Form Processes with Hierarchies	/ 81
2.6.3.3	Form Processes with Aggregate Functions	/ 82
2.6.4	Statistical Transformation	/ 82
2.6.5	Normal Forms and Performance Tradeoff	/ 82

- 2.7 When Is a DBMS Relational? / 83
 - 2.7.1 Foundational Rules (Rule 0 and Rule 12) / 84
 - 2.7.2 Structural Rules (Rule 1 and Rule 6) / 85
 - 2.7.3 Integrity Rules (Rule 3 and Rule 10) / 87
 - 2.7.4 Data Manipulation Rules / 88
 - 2.7.5 Data Independence / 89
 - 2.7.6 Logical Data Independence / 90
 - 2.7.7 Distribution Independence / 90
 - 2.7.8 Advantages of Conforming to Strict Relational Criteria / 90
- 2.8 The Future of Relational Systems / 91
 - 2.8.1 Distributed Databases / 91
- 2.9 Summary / 96

3.

OBJECT ORIENTATION

97

- 3.1 Introduction / 97
 - 3.1.1 The History and Evolution of Object Orientation / 98
- 3.2 Introduction to Object Orientation / 99
 - 3.2.1 What Is Object Orientation? / 100
 - 3.2.2 Model of Reality / 102
 - 3.2.3 Approaches to Object-Oriented Programming / 104
- 3.3 Abstract Data Types / 105
 - 3.3.1 Classes / 109
 - 3.3.1.1 Instance Variables / 111
 - 3.3.1.2 Object Creation / 113
 - 3.3.1.3 Methods and Messages / 114
 - 3.3.2 Dynamic Binding / 116
 - 3.3.3 Constraints / 118
- 3.4 Inheritance Hierarchies / 119
 - 3.4.1 Inheriting Instance Variables / 121
 - 3.4.2 Inheriting Methods / 125
 - 3.4.3 Multiple Inheritance / 126
 - 3.4.4 Metaclasses / 129
 - 3.4.5 Prototype Systems and Delegation / 133
- 3.5 Object Identity / 135
 - 3.5.1 Identity through Identifier Keys / 137
 - 3.5.2 Preserving Identity through Object Modifications / 142
- 3.6 Object-Oriented Databases / 143
 - 3.6.1 GemStone / 145
 - 3.6.1.1 Querying in OPAL / 145
 - 3.6.1.2 Indexing in GemStone / 148

3.6.1.3	Transactions in GemStone	/ 150
3.6.1.4	Clustering in GemStone	/ 151
3.7	Applications of Object Orientation	/ 152
3.7.1	Discrete-Event Simulation	/ 153
3.7.2	Design Engineering	/ 156
3.7.3	Software Engineering	/ 158
3.8	Summary	/ 159

4.

EXPERT SYSTEMS

161

4.1	Introduction	/ 161
4.1.1	Introduction to Expert Systems	/ 162
4.1.2	The History and Evolution of Expert Systems	/ 163
4.1.3	Expert Systems as a New Programming Paradigm	/ 164
4.1.4	Expert System Development	/ 166
4.1.5	The Architecture of an Expert System	/ 167
4.2	Knowledge and Inference	/ 168
4.2.1	Facts and Rules	/ 169
4.2.2	Inference	/ 172
4.2.3	Backward Chaining Inference	/ 174
4.2.4	Forward Chaining Inference	/ 177
4.2.5	Inexact Inference	/ 179
4.3	Objects and Frames	/ 183
4.3.1	Storing and Retrieving Knowledge in Objects	/ 186
4.3.2	Attached Predicates	/ 188
4.3.3	Inheritance	/ 190
4.4	Intuitive Model of Knowledge	/ 192
4.4.1	Components of Knowledge	/ 194
4.4.2	Naming, Describing, and Organizing	/ 197
4.4.3	Relating and Constraining	/ 198
4.4.4	Blackboards	/ 199
4.5	Integrating Expert Systems with Databases	/ 201
4.5.1	Relational Objects and Relational Predicates	/ 203
4.5.2	Defining Schemas	/ 204
4.5.2.1	Defining Indices	/ 205
4.5.3	Dynamic Schema Modifications and Transactions	/ 205
4.5.4	Adding and Removing Data	/ 206
4.5.5	Retrieving Data	/ 207
4.6	Using the Integrated Expert Database System	/ 210
4.7	Dynamic Data Management with Objects	/ 217
4.7.1	Integrity Maintenance with Attached Procedures	/ 217
4.7.2	Dynamic Data Management with Attached Predicates	/ 219
4.8	Summary	/ 221

5.

HYPERMEDIA

223

- 5.1 Introduction / 223
 - 5.1.1 Hypertext: A Definition / 224
 - 5.1.2 Why Hypertext? / 225
- 5.2 The History of the Hypertext Concept / 228
 - 5.2.1 The Memex / 228
 - 5.2.2 NLS / 229
 - 5.2.3 The Xanadu Project / 230
- 5.3 Hypertext Systems / 231
 - 5.3.1 HyperCard / 232
 - 5.3.2 NoteCards / 237
 - 5.3.3 Guide / 239
 - 5.3.4 Other Systems / 240
- 5.4 Knowledge Representation within Hypertext / 242
 - 5.4.1 Node Types / 242
 - 5.4.2 Link Types / 244
 - 5.4.2.1 Navigation Links / 244
 - 5.4.2.2 Organizational and Inferential Links / 245
 - 5.4.3 An Illustration of the Hypertext Model / 246
- 5.5 A Model of Hypertext / 248
 - 5.5.1 Hypertext and Semantic Networks / 248
 - 5.5.2 The Proposed Model / 249
 - 5.5.2.1 Knowledge Representation / 249
 - 5.5.2.2 Linkage to Inference Systems / 252
 - 5.5.2.3 Information Retrieval / 252
 - 5.5.2.4 Browsing / 253
 - 5.5.3 Extending the Model / 254
 - 5.5.4 Alternative Models / 255
- 5.6 Hypermedia Design / 258
 - 5.6.1 Hypermedia and CD-ROM / 259
 - 5.6.2 Hypermedia Design Guidelines / 259
 - 5.6.3 Hypermedia Utilization / 262
- 5.7 Evaluation of Current Hypertext Systems / 262
 - 5.7.1 Strengths / 263
 - 5.7.1.1 Associative Structuring / 263
 - 5.7.1.2 Exploration / 264
 - 5.7.2 Problems / 267
 - 5.7.2.1 Disorientation / 267
 - 5.7.2.2 Training / 268
 - 5.7.2.3 Standards / 269
 - 5.7.3 Browsing Media and Maps / 271
 - 5.7.3.1 The NoteCards Browser / 271

5.7.3.2	The Fisheye Lens Model	/ 273
5.7.3.3	Landmark Views	/ 275
5.7.4	Converting Text to Hypertext	/ 279
5.8	Hypertext Applications	/ 280
5.8.1	Text Processing and Authoring	/ 280
5.8.2	Hypertext and Databases	/ 280
5.8.3	Information Retrieval	/ 282
5.8.3.1	Hypertext and Query Formulation	/ 283
5.8.4	Expert Hypertext	/ 284
5.9	A Hypertext Case Study	/ 285
5.9.1	Project Jefferson	/ 285
5.10	Summary	/ 291

6.

TEXT MANAGEMENT AND RETRIEVAL

293

6.1	Introduction	/ 293
6.2	Information Retrieval	/ 293
6.2.1	Developments in Information Technology	/ 294
6.2.2	The Nature of Information Retrieval	/ 295
6.2.3	A Definition of Information Retrieval	/ 297
6.2.4	Indexing	/ 300
6.2.4.1	Manual Indexing	/ 303
6.2.4.2	Automated Indexing	/ 304
6.2.5	The Search Intermediary	/ 305
6.2.6	Query Formulation	/ 309
6.2.7	Relevance Feedback	/ 311
6.3	Evaluating Retrieval Performance	/ 315
6.3.1	Recall	/ 316
6.3.2	Precision	/ 317
6.3.3	Utilization	/ 317
6.4	Text Management	/ 319
6.4.1	Text Management vs. Information Retrieval	/ 319
6.4.2	The Role of Knowledge	/ 320
6.4.3	Concept Indexing in Text Management	/ 321
6.4.3.1	Concept Indexing with Frame Hierarchies	/ 324
6.4.3.2	Interactive Concept Indexing	/ 326
6.4.4	Text Interpretation	/ 329
6.4.4.1	Templates: The Basic Syntactic Unit	/ 334
6.4.4.2	Semantic Interpretation	/ 335
6.4.4.3	Translating the Templates	/ 337
6.4.4.4	Hierarchical Indexing	/ 339
6.4.4.5	Functions of the Index Hierarchy	/ 342

- 6.4.5 Queries in Text Retrieval / 343
 - 6.4.5.1 Query Contexts / 344
- 6.4.6 Evaluative Summary / 346
- 6.5 Online Databases and Retrieval Services / 346
 - 6.5.1 LEXIS and NEXIS / 347
 - 6.5.2 Medline / 349
 - 6.5.3 SCISEARCH / 349
 - 6.5.4 Dialog / 349
 - 6.5.5 BRS / 350
 - 6.5.6 Dow Jones News/Retrieval Service / 350
- 6.6 The Role of Expert Systems and Inference in Information Retrieval / 351
 - 6.6.1 EP-X: A Knowledge-Based Search Intermediary / 353
 - 6.6.2 Inference on Large Databases / 360
 - 6.6.3 Information Retrieval and Instruction / 361
- 6.7 Forecasting the Future in Information Retrieval / 362
- 6.8 Summary / 365

7.

INTELLIGENT DATABASES

367

- 7.1 Introduction / 367
 - 7.1.1 The Architecture of an Intelligent Database / 368
 - 7.1.2 Requirements for an Intelligent Database / 369
- 7.2 FORM: A Formal Object Representation Model / 372
 - 7.2.1 The Basic Concepts of FORM / 373
 - 7.2.2 Methods and Facts / 375
 - 7.2.3 Parents, Instances, and Inheritance / 376
 - 7.2.4 Presentations and Links / 377
 - 7.2.5 Dialogs: The Building Blocks of User Interfaces / 377
 - 7.2.5.1 Basic Concepts of Dialogs / 379
 - 7.2.6 An Example of FORM / 381
 - 7.2.6.1 Combining Objects / 382
 - 7.2.6.2 Logical Deduction / 382
 - 7.2.6.3 Querying / 383
- 7.3 The Integrated Intelligent Database Engine / 384
 - 7.3.1 Abstract Data Types / 386
 - 7.3.2 Object Identity / 387
 - 7.3.3 Inheritance / 389
 - 7.3.4 Built-in Data Types / 389
 - 7.3.5 Declarative Predicates and Rules / 389
 - 7.3.6 Certainty Factors / 390
 - 7.3.7 Distribution, Concurrency, Resiliency, Versions, and Multiprocessing / 391

7.4	The Intelligent Database User Interface /	391
7.4.1	User-Interface Core Model /	392
7.4.2	Link Types within the Core Model /	394
7.4.3	The Structure of the Core Model /	398
7.4.4	Interface Operations /	401
7.5	High-Level Tools /	403
7.5.1	Knowledge Discovery Tools /	404
7.5.1.1	Statistics and Machine Learning /	406
7.5.1.2	Approaches to Machine Learning /	409
7.5.1.3	IXL (Induction on eXtremely Large databases) /	411
7.5.1.4	Benefits and Applications of Machine Learning /	414
7.5.2	Data Integrity and Quality Control Tools /	415
7.5.3	Hypermedia Management Tools /	416
7.5.4	Data Presentation and Display Tools /	419
7.5.5	Decision Support and Scenario Analysis Tools /	421
7.5.6	Data-Format Management Tools /	423
7.5.7	Intelligent System Design Tools /	425
7.6	Fortune Finder: An Intelligent Database Application /	425
7.6.1	Browsing /	427
7.6.2	The Note Taker /	429
7.6.3	The Photo Feature /	430
7.6.4	Note Organization /	431
7.6.5	Structured Searching /	431
7.6.6	The Questioner and Adviser /	432
7.6.7	Fortune Finder in Retrospect /	432
7.7	Summary /	433

8.

SUMMARY

435

8.1	The Nature of Intelligent Databases /	435
8.2	Overview of Databases /	437
8.3	Object Orientation /	439
8.4	Expert Systems and Databases /	440
8.5	The Role of Hypermedia /	442
8.6	Text Management and Intelligent Databases /	444
8.7	The Intelligent Database Model /	446

BIBLIOGRAPHY

449

INDEX

469

ABOUT THE AUTHORS

479

INTRODUCTION

■ 1.1 INTRODUCTION

This book is about intelligent databases, which can be loosely defined as *databases that manage information in a natural way, making that information easy to store, access, and use*. We use the term *information* rather than *data* here since the intelligent databases discussed in this book deal not only with traditional applications, such as inventory management, but also with knowledge bases, automatic discovery systems, textual data, images, and so on.

The concepts of "ease" and "naturalness" are so well accepted that it is hard to argue with their worth. However, while the definition above is rather general, in Section 1.5 we provide detailed criteria for what makes a database intelligent. These criteria establish that by using intelligent databases users can perform tasks involving large amounts of information that otherwise could not possibly be performed. This relies on defining three levels of database intelligence.

- i. Intelligence of the high-level tools provided by the database.
- ii. Intelligence at the user-interface level.
- iii. Intelligence of the underlying database engine.

These are illustrated in Figure 1.1.

Intelligent databases represent a new technology for information management that has evolved as a result of the integration of traditional approaches to databases with more recent fields such as




Figure 1.1 Three levels of database intelligence.

- Object-oriented programming
- Expert systems
- Hypermedia
- Online information retrieval

The merger of these technologies is illustrated in Figure 1.2.

Intelligent databases thus represent the evolution of a number of distinct paths of technological development. As Figure 1.3 suggests, traditional database systems were developed in the mainframe environments of the early 1960s. They were typically applied to numeric and record-based data (e.g., employee records or parts inventories) and soon began to spread throughout industry. In the early 1970s, online information systems became available as a specialization of database technology, and they began dealing with textual databases. In the mid-1970s the technology of expert systems developed and had become prominent by the late 1970s. Similarly, in the late 1970s and early 1980s, techniques relating to *object-oriented programming* grew out of work in software engineering, user interfaces, and high-level programming languages. Finally, in the mid-1980s practical hypermedia systems, which dealt with information in a variety of forms such as text, images, and sound, became available.

Until recently, these technologies were treated in isolation, with each technology being only *weakly linked* to others. For instance, in the early days expert systems relied on little more than file-transfer protocols to gather data from databases. This was partly