

Zbigniew Michalewicz · Martin Schmidt
Matthew Michalewicz · Constantin Chiriac

Adaptive Business Intelligence

Preface

“My name is Sherlock Holmes. It is my business to know what other people do not know.”

The Adventure of the Blue Carbuncle

“What do you think of it, Watson?”

“A masterpiece. You have never risen to a greater height.”

The Adventure of the Bruce-Partington Plans

Since the computer age dawned on mankind, one of the most important areas in information technology has been that of “decision support.” Today, this area is more important than ever. Working in dynamic and ever-changing environments, modern-day managers are responsible for an assortment of far-reaching decisions: *Should the company increase or decrease its workforce? Enter new markets? Develop new products? Invest in research and development?* The list goes on. But despite the inherent complexity of these issues and the ever-increasing load of information that business managers must deal with, all these decisions boil down to two fundamental questions:

- What is likely to happen in the future?
- What is the best decision right now?

Whether we realize it or not, these two questions pervade our everyday lives – both on a personal and professional level. When driving to work, for instance, we have to make a traffic prediction before we can choose the quickest driving route. At work, we need to predict the demand for our product before we can decide how much to produce. And before investing in a foreign market, we need to predict future exchange rates and economic variables. It seems that regardless of the decision being made or its complexity, we first need to make a prediction of what is likely to happen in the future, and then make the best decision based on that prediction. This fundamental process underpins the basic premise of *Adaptive Business Intelligence*.

Simply put, Adaptive Business Intelligence is the discipline of combining prediction, optimization, and adaptability into a system capable of answering these two fundamental questions: *What is likely to happen in the future?* and *What is the best decision right now?* To build such a system, we first need to understand the methods and techniques that enable prediction, optimization, and adaptability. At first blush, this subject matter is nothing new, as hundreds of books have already been written on business intelligence, data mining and prediction methods, optimization techniques,

and so forth. However, none of these books has explained how to combine these various technologies into a software system that is capable of predicting, optimizing, and adapting. This text is the first on the subject.

When we set out to write *Adaptive Business Intelligence*, we had three important objectives in mind: First of all, we wanted to explain why the future of the business intelligence industry lies in systems that can make decisions, rather than tools that produce detailed reports. As most business managers now realize, there is a world of difference between having good knowledge and detailed reports, and making smart decisions. Michael Kahn, a technology reporter for Reuters in San Francisco, makes a valid point in the January 16, 2006 story entitled “Business intelligence software looks to future”:

“But analysts say applications that actually answer questions rather than just present mounds of data is the key driver of a market set to grow 10 per cent in 2006 or about twice the rate of the business software industry in general.

‘Increasingly you are seeing applications being developed that will result in some sort of action,’ said Brendan Barnacle, an analyst at Pacific Crest Equities. ‘It is a relatively small part now, but it is clearly where the future is. That is the next stage of business intelligence.’”

We could not agree more.

Second, we wanted to explain the principles behind many prediction methods and optimization techniques in simple terms, so that any business manager could grasp them. Even though most business managers have a limited technology background, they should not be intimidated by terms such as “artificial neural networks,” “fuzzy logic,” “evolutionary algorithms,” “ant systems,” or “agent-based modeling.” They should understand the strengths and weaknesses of these methods and techniques, their operating principles, and applicability. Armed with such knowledge, business managers will be in a better position to control the application of these methods and techniques in their respective organizations.

And, third, we wanted to underscore the enormous applicability of Adaptive Business Intelligence to many real-world business problems, ranging from demand forecasting and scheduling, to fraud detection and investment strategies. From a high-level perspective, most of these business problems have similar characteristics, and the application of Adaptive Business Intelligence can provide significant benefits and savings.

To facilitate the discussion in this book, we have divided the chapters into three parts that correspond to the three objectives listed above. In Part I, we present the fundamental ideas behind Adaptive Business Intelligence, and explain the different roles that prediction, optimization, and adaptability play in producing near-optimal decisions. We also discuss the characteristics that many business problems have in common, and why these characteristics increase the complexity of the problem-solving exercise. Furthermore, we introduce a particular distribution problem that is used throughout the text as a running example. Given that the prediction and optimization issues in this example are common to most business problems, it should be relatively easy for the reader to extrapolate this example to many other business domains.

Because countless texts have already been written on the subject of database technologies, data warehousing, online analytical processing, reporting, and the like, we saw little point in rehashing the tools and techniques that are routinely used to access, view, and manipulate organizational data. Instead, Part II of the book discusses the various prediction methods and optimization techniques that can be used to develop an Adaptive Business Intelligence system. The distribution example is continued throughout these chapters, effectively highlighting the strengths and weaknesses of each method and technique. Each chapter in Part II is concluded by a *Recommended Reading* section that provides suggestions for readers who want to learn more about particular methods or techniques.

Part III begins with a chapter on hybrid systems and adaptability, explaining how to “combine” the various methods and techniques discussed in Part II, and how the component of adaptability can be added to the final design. In the remaining chapters of the book, we discuss the definitive solution to the distribution problem that was used throughout the text, as well as the application of Adaptive Business Intelligence to several other complex business problems.

Without a doubt, we believe that business managers of all levels would benefit from this text. Anyone who makes operational and strategic decisions – whether on the factory floor or in the boardroom – will find this book invaluable for understanding the science and technology behind better predictions and decisions. We hope that readers will enjoy reading the book as much as we enjoyed writing it, and that they will profit from it.

We would like to thank everyone who made this book possible, and who took the time to share their thoughts and comments on the subject of Adaptive Business Intelligence. In particular, we would like to express our gratitude to SolveIT Software’s scientific advisory board, which includes Hussein Abbass, Valerio Aimale, Jürgen Branke, Mike Brooks, Carlos Coello, Ernesto Costa, Kalyanmoy Deb, Gusz Eiben, Fred Glover, Philip Hingston, Jong-Hwan Kim, Bob McKay, Marek Michalewicz, Masoud Mohammadian, Pablo Moscato, Michael Rumsewicz, Marc Schoenauer, Alice Smith, Russel Stonier, Lyndon While, Xin Yao, and Jacek Zurada. Our special appreciation also goes to two anonymous reviewers who provided us with many insights and useful suggestions, and to Ronan Nugent, who did a wonderful job of editing this book and helping us make the entire project a success.

Lastly, we would like to thank the most famous fictional detective of all time, Sherlock Holmes, for providing us with the entertaining quotes at the beginning of each chapter. Mr. Holmes remains one of the most famous problem-solvers of all time, and his methodology is based on prediction (“*It is a capital mistake to theorize before you have all the evidence*”), optimization (“*... I had best proceed on my own lines, and then clear the whole matter up once and for all*”), and adaptability (“*I have devised seven separate explanations ... But which of these is correct can only be determined by the fresh information, which we shall no doubt find waiting for us*”). Needless to say, his methodology bears a striking resemblance to Adaptive Business Intelligence! Enjoy.

Contents

Part I: Complex Business Problems

1	Introduction	3
2	Characteristics of Complex Business Problems	9
2.1	Number of Possible Solutions.....	10
2.2	Time-Changing Environment	12
2.3	Problem-Specific Constraints	13
2.4	Multi-objective Problems	14
2.5	Modeling the Problem	16
2.6	A Real-World Example	19
3	An Extended Example: Car Distribution	25
3.1	Basic Terminology	25
3.2	Off-lease Cars	27
3.3	The Problem	28
3.4	Transportation	30
3.5	Volume Effect.....	32
3.6	Price Depreciation and Inventory.....	33
3.7	Dynamic Market Changes	33
3.8	The Solution	34
4	Adaptive Business Intelligence	37
4.1	Data Mining.....	38
4.2	Prediction.....	41
4.3	Optimization	43
4.4	Adaptability	44
4.5	The Structure of an Adaptive Business Intelligence System.....	45

Part II: Prediction and Optimization

5	Prediction Methods and Models	49
5.1	Data Preparation.....	51
5.2	Different Prediction Methods	56
5.2.1	Mathematical Methods	56
5.2.2	Distance Methods.....	62

5.2.3	Logic Methods	64
5.2.4	Modern Heuristic Methods	68
5.2.5	Additional Considerations	69
5.3	Evaluation of Models	69
5.4	Recommended Reading.....	74
6	Modern Optimization Techniques.....	75
6.1	Overview	75
6.2	Local Optimization Techniques	82
6.3	Stochastic Hill Climber	87
6.4	Simulated Annealing.....	90
6.5	Tabu Search	96
6.6	Evolutionary Algorithms.....	101
6.7	Constraint Handling	108
6.8	Additional Issues.....	112
6.9	Recommended Reading.....	114
7	Fuzzy Logic	117
7.1	Overview	119
7.2	Fuzzifier	119
7.3	Inference System.....	123
7.4	Defuzzifier.....	127
7.5	Tuning the Membership Functions and Rule Base.....	128
7.6	Recommended Reading.....	129
8	Artificial Neural Networks	131
8.1	Overview	132
8.2	Node Input and Output	134
8.3	Different Types of Networks	136
8.3.1	Feed-Forward Neural Networks.....	137
8.3.2	Recurrent Neural Networks	140
8.4	Learning Methods	142
8.4.1	Supervised Learning.....	142
8.4.2	Unsupervised Learning.....	146
8.5	Data Representation	147
8.6	Recommended Reading.....	148
9	Other Methods and Techniques.....	151
9.1	Genetic Programming.....	151
9.2	Ant Systems and Swarm Intelligence.....	158
9.3	Agent-Based Modeling.....	163
9.4	Co-evolution	169
9.5	Recommended Reading.....	173

Part III: Adaptive Business Intelligence

10	Hybrid Systems and Adaptability.....	177
10.1	Hybrid Systems for Prediction	178
10.2	Hybrid Systems for Optimization	183
10.3	Adaptability	187
11	Car Distribution System	191
11.1	Overview	192
11.2	Graphical User Interface.....	194
11.2.1	Constraint Handling	195
11.2.2	Reporting	201
11.3	Prediction Module.....	203
11.4	Optimization Module	206
11.5	Adaptability Module	208
11.6	Validation	211
12	Applying Adaptive Business Intelligence.....	215
12.1	Marketing Campaigns	215
12.2	Manufacturing.....	221
12.3	Investment Strategies	224
12.4	Emergency Response Services.....	228
12.5	Credit Card Fraud.....	232
13	Conclusion.....	239
	Index	243

Part I:
Complex Business Problems

1 Introduction

“Like all Holmes’s reasoning, the thing seemed simplicity itself when it was once explained.”

The Stock-Broker’s Clerk

“Before turning to ... matter which present the greatest difficulties, let the inquirer begin by mastering more elementary problems.”

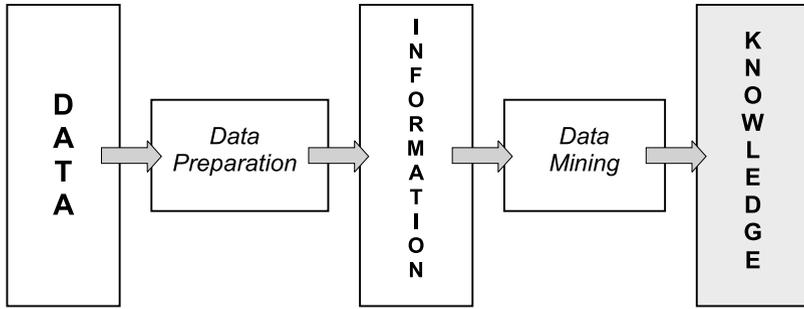
A Study in Scarlet

“*The answer to my problem is hidden in my data... but I cannot dig it up!*” This popular statement has been around for years as business managers gathered and stored massive amounts of data in the belief that they contain some valuable insight. But business managers eventually discovered that raw data are rarely of any benefit, and that their real value depends on an organization’s ability to analyze them. Hence, the need emerged for software systems capable of retrieving, summarizing, and interpreting data for end-users.

This need fueled the emergence of hundreds of *business intelligence* companies that specialized in providing software systems and services for extracting *knowledge* from raw data. These software systems would analyze a company’s operational data and provide knowledge in the form of tables, graphs, pies, charts, and other statistics. For example, a business intelligence report may state that 57% of customers are between the ages of 40 and 50, or that product X sells much better in Florida than in Georgia.¹

Consequently, the general goal of most business intelligence systems was to: (1) access data from a variety of different sources; (2) transform these data into information, and then into knowledge; and (3) provide an easy-to-use graphical interface to display this knowledge. In other words, a business intelligence system was responsible for collecting and digesting data, and presenting knowledge in a friendly way (thus enhancing the end-user’s ability to make good decisions). The following diagram illustrates the processes that underpin a traditional business intelligence system:

¹ Note that *business intelligence* can be defined both as a “state” (a report that contains knowledge) and a “process” (software responsible for converting data into knowledge).



Although different texts have illustrated the relationship between data and knowledge in different ways, the distinction between data, information, and knowledge is quite clear:

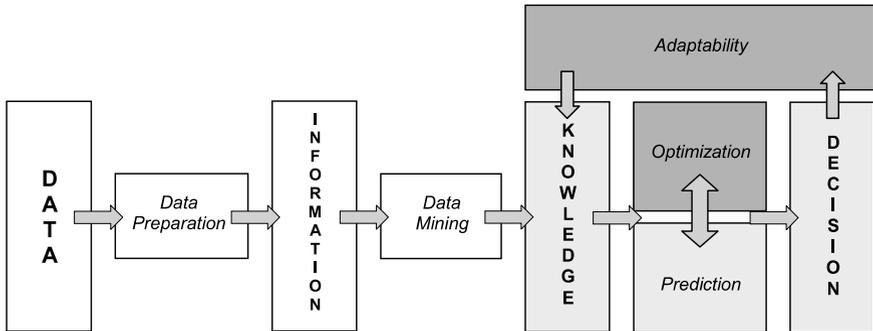
- *Data* are collected on a daily basis in the form of bits, numbers, symbols, and “objects.”
- *Information* is “organized data,” which are preprocessed, cleaned, arranged into structures, and stripped of redundancy.
- *Knowledge* is “integrated information,” which includes facts and relationships that have been perceived, discovered, or learned.

Because knowledge is such an essential component of any decision-making process (as the old saying goes, “*Knowledge is power!*”), many businesses have viewed knowledge as the final objective. But it seems that knowledge is no longer enough. A business may “know” a lot about its customers – it may have hundreds of charts and graphs that organize its customers by age, preferences, geographical location, and sales history – but management may still be unsure of what decision to make! And here lies the difference between “decision support” and “decision making”: all the knowledge in the world will not guarantee the right or best decision.

Moreover, recent research in psychology indicates that widely held beliefs can actually hamper the decision-making process. For example, common beliefs like “the more knowledge we have, the better our decisions will be,” or “we can distinguish between useful and irrelevant knowledge,” are not supported by empirical evidence. Having more knowledge merely increases our confidence, but it does not improve the accuracy of our decisions. Similarly, people supplied with “good” and “bad” knowledge often have trouble distinguishing between the two, proving that irrelevant knowledge decreases our decision-making effectiveness.

Today, most business managers realize that a gap exists between having the right knowledge and making the right decision. Because this gap affects management’s ability to answer fundamental business questions (such as “What should be done to increase profits? Reduce costs? Or increase market share?”), the future of business intelligence lies in systems that can provide answers and recommendations, rather than mounds of knowledge in the form of reports. *The future of business intelligence lies in systems that can make decisions!* As a result, there is a new trend emerging in the marketplace called *Adaptive Business Intelligence*.

In addition to performing the role of traditional business intelligence (transforming data into knowledge), Adaptive Business Intelligence also includes the decision-making process, which is based on prediction and optimization:



While *business intelligence* is often defined as “a broad category of application programs and technologies for gathering, storing, analyzing, and providing access to data,” the term *Adaptive Business Intelligence* can be defined as “the discipline of using prediction and optimization techniques to build self-learning ‘decisioning’ systems” (as the above diagram shows). Adaptive Business Intelligence systems include elements of data mining, predictive modeling, forecasting, optimization, and adaptability, and are used by business managers to make better decisions.

This relatively new approach to business intelligence is capable of recommending the best course of action (based on past data), but it does so in a very special way: An Adaptive Business Intelligence system incorporates prediction and optimization modules to recommend near-optimal decisions, and an “adaptability module” for improving future recommendations. Such systems can help business managers make decisions that increase efficiency, productivity, and competitiveness. Furthermore, the importance of *adaptability* cannot be overemphasized. After all, what is the point of using a software system that produces sub par schedules, inaccurate demand forecasts, and inferior logistic plans, time after time? Would it not be wonderful to use a software system that could *adapt* to changes in the marketplace? A software system that could *improve* with time?

The concept of adaptability is certainly gaining popularity, and not just in the software sector. Adaptability has already been introduced in everything from automatic car transmissions (which adapt their gear-change patterns to a driver’s driving style), to running shoes (which adapt their cushioning level to a runner’s size and stride), to Internet search engines (which adapt their search results to a user’s preferences and prior search history). These products are very appealing for individual consumers, because, despite their mass production, they are capable of adapting to the preferences of each unique owner after some period of time.

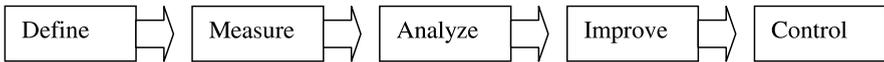
The growing popularity of adaptability is also underscored by a recent publication of the US Department of Defense. This lists 19 important research topics for the next decade and many of them include the term “adaptive”: *Adaptive* Coordi-

nated Control in the Multi-agent 3D Dynamic Battlefield, Control for *Adaptive* and Cooperative Systems, *Adaptive* System Interoperability, *Adaptive* Materials for Energy-Absorbing Structures, and Complex *Adaptive* Networks for Cooperative Control.

For sure, adaptability is here to stay. It is a vital component of any intelligent system, as it is hard to argue that a system is “intelligent” if it does not have the capacity to adapt. Moreover, modern definitions of natural and artificial intelligence include the term “adaptive.” For humans, the importance of adaptability is obvious: our ability to adapt was a key element in the evolutionary process. In the case of artificial intelligence, consider a chess program capable of beating the world chess master: Should we call this program intelligent? Probably not. We can attribute the program’s performance to its ability to evaluate the current board situation against a multitude of possible “future boards” before selecting the best move. However, because the program cannot learn or adapt to new rules, the program will lose its effectiveness if the rules of the game are changed or modified. Consequently, because the program is incapable of learning or adapting to new rules, the program is not intelligent.

The same holds true for any expert system. No one questions the usefulness of expert systems in some environments (which are usually well defined and static), but expert systems that are incapable of learning and adapting should not be called “intelligent”! Some expert knowledge was programmed in, that is all.

It is not surprising that the fundamental components of Adaptive Business Intelligence are already emerging in other areas of business. For example, the *Six Sigma* methodology is a great example of a well-structured, data-driven methodology for eliminating defects, waste, and quality-control problems in many industries. This methodology recommends the following sequence of steps:



Note that the above sequence is very close “in spirit” to part of the previous diagram, as it describes (in more detail) the adaptability control loop. Clearly, we have to “measure,” “analyze,” and “improve,” as we operate in a dynamic environment, so the process of improvement is continuous. The SAS Institute proposes another methodology, which is more oriented towards data mining activities. Their methodology recommends the following sequence of steps:



Again, note that the above sequence is very close to another part of our diagram, as it describes (in more detail) the transformation from data to knowledge. It is not surprising that businesses are placing considerable emphasis on these areas, because better decisions usually translate into better financial performance. And

better financial performance is what Adaptive Business Intelligence is all about. Systems based on Adaptive Business Intelligence aim at solving real-world business problems that have complex constraints, are set in time-changing environments, have several (possibly conflicting) objectives, and where the number of possible solutions is too large to enumerate. Solving these problems requires a system that incorporates modules for prediction, optimization, and adaptability. In the following chapters of this book, we will discuss these modules in detail, and see how they are combined to create an Adaptive Business Intelligence system.