Introduction

The Maturing of AI:

From Science to Technology and Back

Alain Rappaport

Artificial intelligence is a science or a branch of computer science. But it is now also becoming a technology for the real world. This event is happening at the right time, when the complexity of today's business environment is increasing faster than human skill. In all industries and government bodies, events happen at a pace requiring the information infrastructure to quickly adapt (figure 1). This adaptation cannot occur using conventional techniques.

Until recently, modern information systems could withstand the demand of the user by somehow bringing more power than could be used, more languages and tools than could be learned. The remarkably rapid adoption of these new techniques, from computer-aided design to databases and personal computing software, has generated complex environments that require a new level of computing. The infrastructures need to be managed, the software needs to be glued, the productivity of application development needs to be increased, and the development cycles need to be reduced. It is the combination of the systems' complexity and the fluidity of today's business contingencies

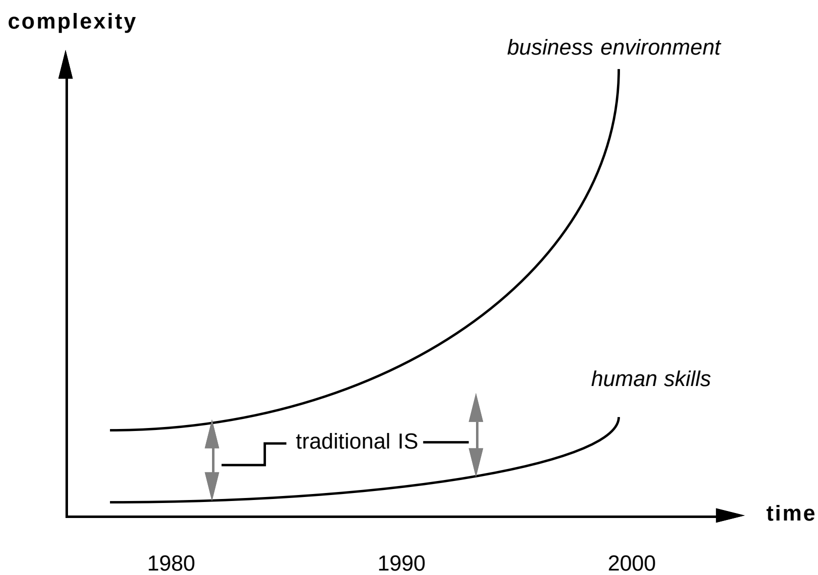

Figure 1. Evolving Complexity of the Business Environments (IS = Information Systems).

(that these systems help establish) that requires a new generation of software. Human skills simply cannot evolve fast enough and need to be better used.

It is clear that the environment we create evolves far faster than these limited resources. Undoubtedly, we need to either manage this environment or have it manage itself. This feat can only be achieved by manipulating knowledge. It is not just a research perspective for some future world; computing with knowledge has become a natural need, arising from the complexity of the world we built.

Many examples illustrate that this knowledge computation is becoming a reality. AI has emerged as the technology that allows us to bridge the gap, adapt our existing environment to the rapidly changing conditions, and create new software applications designed to be an extension of not only the physical body but also the mind. As shown in figure 2, the place of this new technology is on top of the existing one. One result of this interdependence between conventional and AI technology is the increased attention being paid to integration of AI and conventional technology. Just as the mind without the body is meaningless, there also cannot be an "AI-centric" view of the world.

Some examples of this evolving complexity follow. If these global

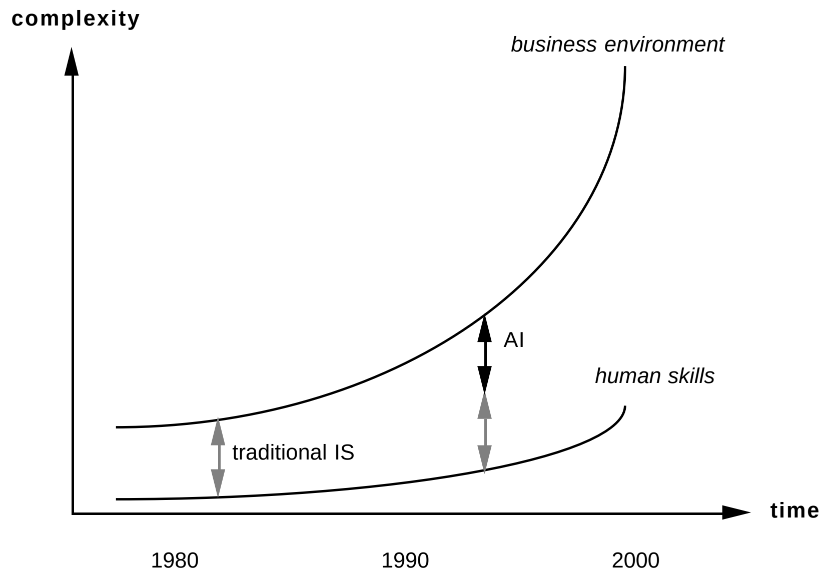

Figure 2

issues are not currently being addressed by AI technology, they are obvious candidates.

The first example is competitiveness, which requires tools for better, more enhanced productivity. For example, software design is a critical aspect of high-technology industrial design, where the shape of the product is part of the product. There are also applications in design for manufacturability, expressing manufacturing constraints at the design level and reducing the manufacturing design cycle.

Second is deregulation. Deregulation increases the degrees of freedom and the complexity of many businesses. In finance, the automatic monitoring of foreign exchange transactions to ensure they are within the correct fiscal, legal, and other boundaries is a new type of application.

Third is regulation, the natural counterpart to deregulation. The evolution of the health system or of environmental concerns makes dealing with regulation an important issue for the use of AI. Related problems exist in areas such as insurance, finance, and telecommunications.

Fourth is geopolitics. Political, economic, and military contingencies are also rapidly changing, affecting policies; changing the directions of programs toward monitoring and verification, for example; and reassessing local and global issues. AI is already being used to measure the effects of policy changes. The pace of the modern world also leads to a general overload of information for the decision maker. Classifying telexes or selecting news stories are AI applications that directly help in dealing with vast amounts of information, bringing it to a manageable level for decision makers.

* * *

It is important to convey the actual emergence of this technology as a major component in a broad spectrum of activities, as an invasive process demonstrating and increasing the participation of AI in our information infrastructure. In this book we see several examples of AI technology being applied to a variety of domains:

First is foreign exchange monitoring. The novelty of the Inspector system lies in both its implementation and its professional implications ("Inspector: An Expert System for Monitoring Worldwide Trading Activities in Foreign Exchange," p. 16). It is an embedded system that performs a safeguard role in the background of international trading activities that involve transactions of varying sizes. This system illustrates the integration of AI techniques with other standard technologies. The payoffs are unusual: They are mainly preventive, but a single discovery is likely to be worth many times the investment made to build the system.

The second domain is software validation. Comparing text output using regression testing programs is the area of the intelligent comparison tool ("Intelligent Text Comparison in Software Validation," p. 246). Instead of being left to the expert, the comparison criteria are now embedded in the comparison software, enhancing productivity. This application illustrates the increasingly important role of AI techniques that are embedded in software-engineering tools.

Third is microfossil identification. This domain is certainly new to most people; it is well known, however, to the oil industry, where it is important for oil exploration. The Vides system helps to visually and rapidly identify what type of microfossil is present at a drilling site ("Integrating Artificial Intelligence and Graphics in a Tool for Microfossil Identification for Use in the Petroleum Industry," p. 202). It is a good illustration of how specialized and narrow scientific domains can be of critical importance in today's competitive world.

These three applications illustrate the use of AI to reduce the complexity of a domain or task. In the following examples, AI is obviously making its way in areas more likely to affect some aspect of our daily life.

First is human resources. The Resumix system allows the intelligent dispatching of resumes to persons within a firm who have needs for employees with certain qualifications ("Computers Assist Humans in Human Resources," p. 179). Human resources and hiring are critical to companies. The use of search, optical character recognition systems, and other office automation tools is again indicative of the role that AI systems are taking as components in existing or new computing infrastructures.

The second domain is news story indexing. Streams of news, indicators, and other numbers or images are part of the daily business environment. The volume of electronic text continues to increase, and it is critical for those using such information that it is correctly classified and indexed ("Construe-TIS: A System for Content-Based Indexing of a Database of News Stories," p. 49). Online systems such as Construe-TIS are used by an increasing number of professionals who rely on rapid access to information, from general data to current global events and trends, to make their decisions. Without such systems, these professionals have no way to manage the voluminous amount of available information.

Third is telex classification. Other techniques are now being used to help in this general problem-solving framework. CASE-based reasoning was efficiently put to use in the PRISM system to classify telexes according to their content ("PRISM: A CASE-Based Telex Classifier," p. 25). This example again illustrates the use of AI as a technology to manage interaction and the flow of information.

Concerning the addition of AI to a large spectrum of existing technologies and products, there is little doubt that many of these technologies, such as databases, document processing systems, or computer-aided design systems, are underused; they need to be adapted to the mental model of the user or the organization to be efficient. The examples previously mentioned reflect a trend toward meeting this need. They also reflect the close integration of AI with existing products from established traditional technologies. There is no doubt that the needs and the expectations we have as citizens, individual consumers, or organizations have shifted toward a customization of the products we want to use. That is, the need exists to have things function the way we want, almost with the desire to design the products ourselves, be they hardware or software. In this book, we see how AI can help.

We are witnessing a rapid extension of some basic definitions. Only a few years ago, hardware meant chips for computing, and software meant lines of code. Hardware now covers a much broader spectrum, from computers to video cameras to high-density television broadcasting to smart cards and teller machines. Software now covers code sound, movies, and knowledge bases as well. AI will play a major role in the integration and expansion of these technologies. We are already witnessing a transition in AI from traditional problem solving, where the problem is predefined, to a looser (no less complex, however) problem solving, where an extension of the mind, from the perceptual to the executive levels, is needed to carry out activities. Systems can have these new resources, extending their capabilities to apply chunks of knowledge, adapt to new situations, or perceive new patterns.

AI has become a highly promising area. Different AI paradigms and architectures complement the expert system approach. New approaches are based on integrating different AI techniques ("Cooperating Artificial Neural and Knowledge-Based Systems in a Truck Fleet Brake-Balance Application," p. 261). In general, one should expect important developments in the areas of machine learning and knowledge acquisition. In machine learning, it has become important to design with operationalism in mind. In knowledge acquisition, it has become critical to build practical techniques that can stand the trial of real development. We should soon see significant applications that involve machine vision and speech as well as program synthesis.

What is changing is the definition of progress in the field. It now includes real-world testing and integration. Fundamental research results need to undergo a software-engineering process that turns these results into technologies that can be tested within real environments, much like the process of clinically testing a drug. The clinical trial of AI is proving most successful, which is why we can speak of going from science to technology and back. Examining the designs, results, and effects of real-world applications is fundamental to the good design of future AI software technologies. This feedback is important to scientific progress, just as it is in other mature scientific and technological fields. It is time to close the loop from science to technology and from the application back to science.

We are still in the interesting zone where different concepts and perspectives are tested while important applications are developed, fielded, and assessed. Some applications, demonstrated in the lab, are still too hard to replicate in the field, while experience gained in fielding others paves the way for new research. There are good indications that it will be difficult to design future software without embedding this new technology in them.

To doubt that AI, in all its different facets, is rapidly becoming a basic technology that is critical to addressing the complexity we have generated and that must efficiently be managed is, in my view, to speak against the evidence.