

From Information to Knowledge: The continuous need to teach AI

Abstract. The paper discusses the importance of teaching Artificial Intelligence as part of a computer science curriculum. The BRETAM model is used to study the advancement from computer systems to computing systems and it is further used to analyze the historical development of AI, as well as to put in perspective the contributions of AI to computing. Finally the paper discusses what may constitute an appealing and influential AI undergraduate course.

1 Introduction

The AI community celebrates this summer its fiftieth year. It was in 1956 during the Dartmouth Summer Research Conference that the term Artificial Intelligence was coined by its founders (see [1] p. 4).

In a more local context, the Mexican AI community celebrates the twentieth anniversary of the foundation of the Mexican Society of AI (SMIA). After three successful yearly workshops on AI research in Mexico José Negrete [2] and a group of young AI Mexican researchers decided to organize the SMIA. The first workshop had been held in Guanajuato in the summer of 1984 with as few as eleven papers [3]. The second meeting was held in Monterrey and the third one in Oaxaca in 1986. It was shortly after that with 25 members in August 1986, that the SMIA was reluctantly founded.¹ From its modest beginnings, the SMIA grew into a full organization that has promoted national and international conferences such as the numerous editions of the RNIA, MICAI and IBERAMIA [4] as well as the establishment of the Mexican Society of Computer science (SMCC).

It is interesting to note that, as such, the Mexican AI community had been publishing in different conferences without having a formal society [5]. The case was similar in the US, although AI began in 1956, the AAAI was only chartered in 1979. Similarly, the theory and practice of AI has been part of the mainstream of computing although not always being properly recognized as AI. Notwithstanding this opacity, teaching of AI and PhDs in AI, have been part of the computer science curriculum since the very beginning of the field (see [6] p.256). In light of the two anniversaries we are celebrating and that ambivalent presence of the discipline within Computer Sciences,

¹ Weary of past experiences, the founders of SMIA decided to sub denominate it “a society of friends” to make explicit its ultimate character and purpose while capitulating to the objective need of having an institutional entity to represent Mexico in academic meetings with societies from various countries, specifically with the Spanish and Portuguese AI societies. Notably, José Cuena, president of the Spanish AI society (AEPIA), had by then managed to convince Spanish and Mexican communities to set up a cooperation agreement on AI whose signature required formal counterparts.

it may not seem out of place to revisit the role of AI in the computing sciences curricula.

Specifically, in this paper we argue that it is wise to teach AI as part of the core computing curriculum and we further propose the format of a core undergraduate AI curriculum. We base our proposal on the historical contributions of AI to computing and support that approach with the BRETAM model which has thus far been used for technological forecasting [7]. We also show how our proposal matches against the standard AI curricula. Although the paper might be read in this forum as a preaching to the faithful, there is a novel argument for C. Sc. curriculum design that pinpoints the deep influence of AI in the IT field and proposes a core curriculum for AI formulated through an explanatory model --predicated on historical and dynamic components-- that has not been explicitly used for these purposes before. To cover these matters, the paper discusses the role of AI in the IT curricula (Section 2), a quick view of the BRETAM model (Section 3) and its application to AI breakthroughs (Section 3).

2. AI in the Computer Science Curricula

*"Computer science is no more about computers
than astronomy is about telescopes." Edsger Dijkstra*

Today Computer Science would be much better characterized if it were called *Computing Sciences* since it is more related with the study of logic, systems and the relationship with discrete mathematics than with hardware itself. In this same spirit, although the ACM was initially the Association for Computer Machinery, it changed its name to Association for Computing Machinery long time ago. Likewise the curricula in most IT programs today reflect this emphasis [8]. However that has not always been the case. To understand this switch in emphasis a frame of reference might be useful. For this purpose we have chosen the BRETAM model proposed by Brian Gaines and Mildred Shaw [7]. We will delve more into it later in the paper, but for the moment let us mention that the main idea of the model considers technological breakthroughs as the guiding principles for technological innovation and development, and focusing in AI innovations we shall apply the model to AI curriculum design.

Even in the early days of computer science, when hardware and basic programming were the major concerns of computer users, the presence of AI was considerable in defining the path for algorithms and data structures. One needs only remember IPL and LISP as the first examples of pointer based programming. Concepts like Data Structures and Query languages grew out of AI research as did the relationship between problem solving and search. It is interesting to read the introduction of one of the first books on Data Structures by Horowitz & Sahni [9] where AI research is mentioned as the common place for student to learn about data structures and algorithms. Indeed many of such algorithms were initially developed to create data representation for both information and knowledge. However, the notions of data

representation and search algorithms came from AI with a concept that was new and became fundamental to computer science: knowledge processing. From its very beginning, rather than just dealing with algorithms, AI was also concerned with problem solving and providing solutions based on strategy, experience and common sense.

The central notion being discussed in those early years of computer science was what Simon [1] called “complex information processing” or more technically, the exploration of the realization that computers were in fact symbol manipulators. The direct outcome of that research question was the development of data structures that were practical for dealing with symbolic problems: mainly lists and trees and languages for concurrent list manipulation. These developments set the theoretical and technological grounds for the most fundamental topics of computer science: data structures, algorithms and programming languages. Likewise the AI proposal that computers could be able to solve problems that require intelligence was the foundation for the whole topic of search and heuristic algorithms. In the case of both notions, symbolic manipulation and problem solving, the outcome of AI research and the ensuing computer science developments gave rise to significant social artifacts. In particular, of the many resulting social artifacts, the metaphor of *mind as an information processor*, which is at the heart of cognitive psychology, has become part of the intellectual heritage of the twentieth century. Other, more earthly developments, like time-sharing and concurrent computation are also outcomes of these pioneering efforts whose AI ancestry may be obscured by their ubiquity and their pervasive influence. A quick review of other ill defined problems that caught the attention of AI researchers, show similar innovative outcomes, mostly in terms of fertile ideas of knowledge construction using heuristics rather than a the search for an algorithmic solution.

Although AI approaches have been polemical because they question the fundamentals of conventional science, it is through innovative breakthroughs that scientific and technological products become part of our everyday reality. In fact the roots of the word innovation stands for both the novelty (i.e. novation) and the creation of new routine or ways to do things (i.e. in). And in this way AI has contributed to do things differently. In Table 1 we outline some of the more significant contributions. As we shall discuss below, it is feasible to recast those contributions as breakthroughs and then organize the AI-specific curriculum around them using Gaines’s BRETAM model.

Table 1. AI in Computing and IT applications

AI Notion	CS topic	AI-derived technologies
Information Processing	List languages, list and tree data structures, concurrent interactions	Mind-as-computer metaphor, time-sharing, general purpose programming languages.
Problem Solving	Search, decision-making	Chess machines, search engines
Human Computer Interaction	Natural language processing, scene recognition, robotics	Pattern matching, image segmentation
Expertise	Knowledge vs. information, automated deduction, Neural networks, pattern recognition, case-based-reasoning	Knowledge based systems
Learning	Genetic algorithms, evolutionary computing	Biometric devices, stock market forecasting
Evolution and adaptation	Agents, BDI architectures	Design optimization, scheduling, P-families of NP problems
Autonomy	Open systems, MAS	NBIC devices, ubiquitous and embedded computing
Coordination		Semantic web, GRID computing, reconfigurable networks

The rationale for keeping AI contents as part of the IT curricula is grounded in the historical fact that AI has been systematically contributing to mainstream computing and providing technologies that are substantial elements of the IT landscape. That significance was present in the original program of AI as stated in the proposal –by McCarthy, Minsky, Rochester and Shannon-- for the Dartmouth Conference ([10] p. 93):

“We propose that a two-month, ten-man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it”

Although in the Dartmouth proposal the research topics are stated in a very general fashion, we can readily see that they are still familiar to us. The purpose then and now is the development of tools to model data at higher level of abstraction; able enough to write programs that closely resemble what society calls human intelligence. In any case to achieve those goals one needs to model and represent knowledge and skills to prototype the tools needed in AI programs.

Many innovations have taken place since the heroic years of nascent AI. Over these years, the computer science curriculum was established and gone through many changes to reflect the advances and needs of the time. Today Information Technology plays a leading role in the world we live in: a world of information where massive amount of people are fully connected and where almost any piece of information is accessible on-line and may be downloaded for a small price. This reality puts into question many long-held practices in computing. For instance, office automation tools and the use of software packages seem to have taken over the need to learn programming skills, efficient web indexing services seem to overcome the need of

traditional enterprise data documentation practices as much as web-based computing seems to overcome the compatibility barriers imposed by vendor set standards; we seriously question this confident oversimplification of IT. Granting that innovations change the skills and knowledge computing professionals need to master and that the professional profiles of IT professionals become more specialized and diverse, it is no less true, however, that there is a core of knowledge and skills --including programming skills, metadata and semantic labeling and interoperability--that are still needed to practice IT *professionally* in any specialty. Such core contents have been discussed --as we hinted a few paragraphs back-- in terms of the professional skills practitioners should master, and also around the curricula of the main IT specialties.² Among these core contents, concerning theory, abstraction and design AI is present as one of the nine study areas for all the IT specialties proposed by the ACM Core Computing Task Force [12]. In later documents AI is again recognized as one of the core technologies of IT and also through the types of problems, approaches and tools AI has brought to IT in general (Figure 1):

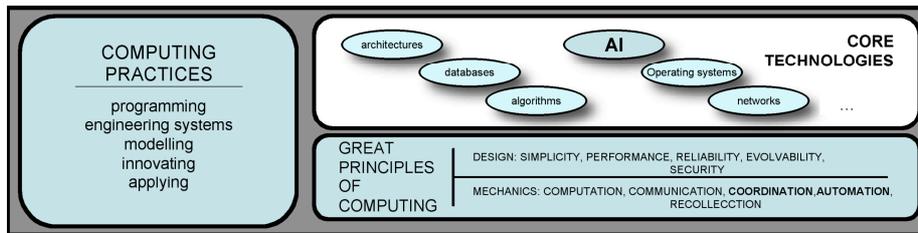


Figure 1. AI is chosen as one core IT technology, also as part of the “*great principles of computing*” and of “*computing practices*” (From Denning [13])

The above comments should lend objective support to the obvious claim that AI has historically been of fundamental importance for the advancement of computing sciences, and if history is in any sense a good indication for the future AI should continue to play that role. We have also shown that AI has been and continues to be a fundamental part of the computing sciences curricula. We believe that it is worth keeping AI in the core IT curricula in order for AI to keep playing its historical positive role and we will suggest the core AI contents to fulfill that purpose.

3 The BRETAM Model

In the introduction we talked about the need to have a framework to explain the advancement and evolution of computer technology and suggested that such framework could also help in designing an AI curriculum. For our discussion we find

² The latest ACM-IEEE Computing Curricula documents [8,11] propose curricular guidelines for five major IT professional profiles: Information Systems, Information Technology, Computer Engineering, Software Engineering and Computer Science.

the BRETAM model particularly suited to satisfy the focus given in this paper. Initially it was proposed in a seminal paper by Gaines and Shaw [7]. The model is based on the claim that technological breakthroughs are the guiding principles for innovation and the further development of technology. Figure 2 presents a depiction of the model.

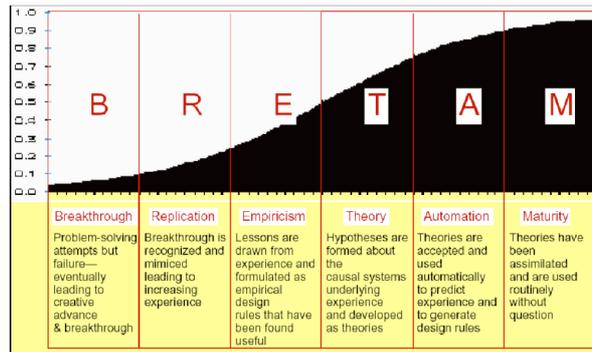


Figure 2. The BRETAM model (taken from Gains and Shaw [7])

The model assumes six phases in the development of a given technology, namely:

1. *Breakthrough* => *A creative new approach or discovery*
2. *Replication* => *Breakthrough is confirmed by new examples and experiences.*
3. *Empiricism* => *Lessons from experience are translated into a somewhat schematic form*
4. *Theory* => *Hypotheses and theories are used to explain findings*
5. *Automation* => *Theories are used to explain experience cast into design mechanisms*
6. *Maturity* => *Assimilation of the technological breakthrough*

As shown in Figure 2, the six phases of the BRETAM model label the well-known logistic “learning curve” in the background (for example see Crane [14]), thus establishing a correspondence between the type of activities involved in developing a technology (introduced in this model) and the productivity of investment as a technology matures (as depicted by the traditional logistic curve). The concept of a technological breakthrough must be stressed here in terms of the creative and innovative aspects involved in the discovery process while researching in a previous paradigm. The issue of knowledge as an insight stemming from experience is very much part of the AI heuristic approach.

An important feature of Gaines’s model is the acknowledgment of tiers or levels of BRETAM phases that occur in a staircase fashion for a sequence of technologies in which each technology is the foundation of the next. The focus of attention in research, product innovation and so on, moves from one technological breakthrough to the next ones. In this way the model proposes a convenient way of identifying the stages of the learning process involved in the maturation of a technology. Innovation in a new technology builds on new discoveries, which are grounded on empirical findings that come from the R phase of its preceding technology. Research in turn

Similarly the books more often in use for this course, namely Winston [15], Russell and Norvig [16], Nilsson [17] and Coppin[18] cover the material extensively. In any case it must be stated that the topic not properly discussed in them are genetic algorithms and that for the case of adaptive learning, neural networks seems to be the primary example used. But overall the various BRETAM tiers are –as expected-- properly covered in the literature. Yet the material is so extensive that the course is either very difficult and arcane or the topics are covered in a rather shallow way.

As a product of our pedagogical experience of more than twenty five years in teaching AI at various institutions, we propose here a course that allow students to comprehend the issues of heuristic functions, learning, selection and agents in an engaging and alluring manner and yet focused enough to provide student a reasonable good discussion of the various paradigms. For this we follow four basic guidelines. First the course is divided in four parts to focus on one specific pedagogical objective for each part:

1. *Symbolic AI => Knowledge representation, search, expert systems*
2. *Learning => Adaptive learning, neural networks, collective learning automata*
3. *Selection => Genetic and evolutionary algorithms*
4. *Agents => Autonomy, agenticity, collaboration and cooperation*

Secondly, we agree with many AI authors such as Seymour Papert [19], Patrick Winston [15], Roger Shanck [20] in “*Learning by Doing*”; the best way to teach a concept, is by constructing knowledge, i.e. students must write their own code. Thirdly, to get students engaged while comparing and contrasting various AI approaches, they are required to solve a similar problem using three different paradigms and then evaluate the benefits of each. A case in point is to program three versions of a complete zero-sum game, in this way they design and program a computer game using: minimax search, learning and genetic algorithms. Finally, the program interface and language selection must be compatible among the three games in order to fully test and compare the benefits of each approach. Note here that students are required to use both aspects of AI, the algorithmic part (e.g. representation and search) and the heuristic (e.g. evaluation functions, learning). We have found this to be a most rewarding experience for the class. We have found this to be a most rewarding experience for the students. A detailed syllabus, a complete set of power point presentations and other course material is available on-line in the following site: (**NOTE of the Authors:** SITE OMITTED FOR BLIND REVIEWING PROCESS).

6 Final Remarks

We have presented the guidelines for a core undergraduate course in AI. It is based on the recognition that there are significant breakthroughs in AI that have influenced the development of IT in a positive way. The course is organized around four tasks that

allow students to get a hands-on idea of these breakthroughs and obtain a general perspective of the main contributions of AI to computing science.

We used the Gaines-Shaw model for technological forecasting to describe AI breakthroughs and correlated them with Gaines's own list of technological breakthroughs in computer and cognitive science. That correspondence indicates the strong connection of AI with mainstream computing. We took advantage of these insights to recast traditional AI contents into our proposed syllabus.

Most importantly we presented evidence of the positive influence of AI in computing by enumerating major AI areas and how they spawned significant IT products. We further inspected the curricula proposals for IT professionals and showed the pertinence and persistency of AI contents in those curricula. If anything, we have presented elements to hold an optimistic view of the role AI is called to play in the future of ITs and therefore on the pertinence and significance of teaching sound undergraduate AI courses in all IT specialties.

References

1. Simon, H.: *The Sciences of the Artificial*. (Third Ed.). MIT Press, MA (1996)
2. Negrete J.: *Inteligencia aunque sea Artificial*. Limusa, Mexico (1990)
3. Lemaitre C., Sanchez, A.: *Las reuniones anuales de la SMIA*. In Galindo Soria, F. (ed) *La Inteligencia Artificial en México*. UTM. Oaxaca (1992)
4. Cantú, F.: *Inteligencia Artificial en México: Historia, estado actual y perspectiva*. In proceedings of TAINA the Workshop on Artificial Intelligence, CIC-IPN, Mexico City (1998)
5. Lemaître, C.: *Inteligencia Artificial en México: Una Visión Personal*. In *Soluciones Avanzadas* (1994), Número 13
6. Simon, H.: *Models of my life*. Basic Books. N.Y. (1991).
7. Gaines, B.R., Shaw, M.L.G.: *A learning model for forecasting the future of information technology*. *Future Computing Systems* 1 (1986) 31-69.
8. ACM, IAS, IEEE: *Computing Curricula 2005: The Overview Report*. ACM SIGCSE Bulletin Volume 38, Issue 1 (March 2006)
9. Horowitz, E., Sahni, S.: *Fundamentals of Data Structures*, Computer Science Press, Rockville, MD (1978).
10. McCurdock, P.: *Machines*. W.H, Freeman and Company. San Francisco (1976).
11. ACM, IEEE: *Computing Curricula 2001*. Joint Task Force on Computing Curricula. IEEE Computer Society, Association for Computing Machinery. *Journal on Educational Resources in Computing (JERIC)*. Volume 1 , Fall 2001
12. Denning, P.J., D.E. Comer, D. Gries, M. C. Mulder, A. Tucker, A.J. Turner and P. Young: *Computing as a Discipline*. The final report of the Task Force on the Core of Computer Science endorsed and approved by the ACM Education Board. *Communications of the ACM* V. 32, N. 1 (January 1989). 9-23.
13. Denning, P. J.: *Great principles of computing*. ACM Press, N.Y *Communications of the ACM*, V. 46, N. 11, November (2003) Vol. 46, No. 11
14. Crane, D.: *Invisible colleges: diffusion of knowledge in scientific communities*. The University of Chicago Press. Chicago (1972)
15. Winston, P. H.: *Artificial Intelligence (Third Edition)* Addison Wesley, MA. (1992)
16. Russell S. and Norvig. P.: *Artificial Intelligence: A Modern Approach*, Prentice Hall (Second Edition) Saddle River, Allyn Bacon NJ (2003).

17. Nilsson, N.: Principles of Artificial Intelligence. Tioga Publishing Co. Tahoe City, Ca. (1980)
18. Coppin, B.: Artificial Intelligence Illuminated, Jones & Barlett, NY (2004)
19. Papert, S.: Mindstorms, Basic Books, Perseus NY, (1993)
20. Shank, R.: Virtual Learning: A Revolutionary Approach to Building a Highly Skilled Workforce. New York, NY: McGraw Hill, (1997)