

A Multi-Institutional Project-Centric Framework for Teaching AI Concepts

Ingrid Russell

University of Hartford, irussell@hartford.edu

Zdravko Markov

Central Connecticut State University, markovz@ccsu.edu

Abstract - The paper presents the second phase of a curriculum project that builds on existing successful work. Our work involves the development, implementation, and testing of an adaptable framework for the presentation of core AI topics that emphasizes the relationship between AI and computer science. Under phase 1, we developed and pilot-tested our proof-of-concept. The second phase involves further development and testing of this adaptable framework for the presentation of core AI topics through a unifying theme of machine learning using a project-based and application-oriented approach. A total of twenty-six projects that span a variety of applications are being developed and tested. Each project involves the development of a machine learning system in a specific application. The applications span a large area including network security, recommender systems, game playing, computational chemistry, robotics, conversational systems, cryptography, web document classification, vision, bioinformatics, pattern recognition, data mining, and others. We present an overview of this work along with preliminary classroom experiences.

Index Terms – Artificial Intelligence Education, Computer Science Curriculum Issues, Project-Based Learning.

INTRODUCTION

An introductory Artificial Intelligence (AI) course provides students with basic knowledge of the theory and practice of AI as a discipline concerned with the methodology and technology for solving problems that are difficult to solve by other means. It is believed by many faculty members that an introductory AI course is challenging to teach because of the diverse and seemingly disconnected topics that are typically covered [5]. Recently, work has been done to address the diversity of topics covered in the course and to create a theme-based approach. Russell and Norvig present an agent-centered approach [12]. A number of faculty have been working to integrate Robotics into the AI course [1, 2, 3, 4].

Our proposed work incorporates machine learning as a unifying theme for the AI course, through a set of hands-on lab projects. Machine learning is inherently connected with the AI core topics and provides methodology and technology to enhance real-world applications within many of these topics. Machine learning also provides a bridge between AI

technology and modern software engineering. As Mitchell [9] points out, machine learning is now considered as a technology for both software development (especially suitable for difficult-to-program applications or for customizing software) and building intelligent software (i.e., a tool for AI programming).

Some courses seek to use robotics to unify diverse AI topics. However, in an introductory course one wishes to impart a wide variety of topics efficiently, indexing the major areas of the field. This requires that an experiential education allow for *rapid prototyping*. Unfortunately, robotics does not allow for rapid prototyping. In effect, the engaging application limits the breadth of such a course. Robotics teaches through hard experience that a simulated world is often much easier to handle than the real one.

By contrast, a machine learning application can be rapidly prototyped, allowing learning to be grounded in engaging experience without limiting the important breadth of an introductory course. Our machine learning emphasis acknowledges that intelligent systems are best taught through their application to challenging problems. However, by removing the interaction with the physical world, we protect students from getting bogged down with gear ratios and hysteresis effects, allowing them to focus on the rapid development of a variety of problem-solving techniques.

Our proposed work will undoubtedly draw comparison to the introductory text by Nils Nilsson, *Artificial Intelligence: A New Synthesis*, which takes an evolutionary, agent-based approach to topic unification with a heavy emphasis on machine learning [10]. Nilsson organizes the material from perceived low cognitive functions to high cognitive functions, starting with reactive agents and concluding with reasoning agents that communicate with each other. The text is geared toward a student with significant prior coursework in mathematics. In contrast, our approach would allow for varying levels of mathematical sophistication, with implementation of concepts being central to the learning process. In the following sections, we present an overview of phase 2, a summary of the projects being developed, and preliminary classroom experiences.

PROJECT MLEXAI PHASE II

In phase 2 of MLeXAI, we build on the success of our smaller scale phase I project. In Phase I, we developed and pilot-tested our proof-of-concept which included the development of six hands-on laboratory projects that were implemented and pilot-tested at three diverse institutions: University of XXX, a comprehensive private university, YYY University, a large public University, and ZZZ College, a small private college.

Guided by results of our experiences with this prior work and by the positive results of its assessment, we have been extending this curricular material [6, 7, 8, 11]. Specifically, this is a multi-institutional effort that engages a community of 20 scholars, including the two PIs, from a broad range of universities working together on the refinement of existing curricular material, and on the development, implementation, and testing of additional material that extends phase 1 work.

A total of twenty-six student projects that span a variety of applications are being developed and are being tested. Each project involves the development of a machine learning system in a specific application. The applications span a large area including network security, recommender systems, game playing, intelligent agents, computational chemistry, robotics, conversational systems, cryptography, web document classification, vision, data integration in databases, bioinformatics, pattern recognition, data mining, and others. The institutions represent diverse users in several settings.

We are producing a laboratory manual consisting of a suite of adaptable, self-contained, hands-on laboratory projects that can be closely integrated into a one-term AI course. An instructor's manual will also be developed that provides guidance for the adoption and adaptation of this curricular material. The instructor manual will also include sample solutions as well as supporting material such as code and documentation. The effectiveness of this project is being evaluated with the assistance of internal and external evaluators through a multi-tier evaluation system involving faculty, students, and an advisory board. The material developed is being tested at 18 institutions and on diverse users in several settings.

We have developed a suite of adaptable, hands-on laboratory projects that can be closely integrated into an introductory AI course. Each project involves the design and implementation of a learning system which will enhance a particular commonly-deployed application. The goal is to enhance the student learning experience in the introductory artificial intelligence course by (1) introducing machine learning elements into the course, (2) implementing a set of unifying machine learning laboratory projects to tie together the core AI topics, and (3) developing, applying, and testing an adaptable framework for the presentation of core AI topics which emphasizes the important relationship between AI and computer science in general, and software development in particular.

GOALS AND OBJECTIVES

The difficulties mentioned earlier associated with the introductory AI course, combined with the increasingly important role of machine learning in computer science in general and software development in particular, are the motivating factors for this project. The specific objectives are listed below:

- Enhance the student learning experience in the AI course by implementing a unifying theme of machine learning to tie together the diverse topics in the AI course.
- Increase student interest and motivation to learn AI by providing a framework for the presentation of the major AI topics that emphasizes the strong connection between AI and computer science.
- Highlight the bridge that machine learning provides between AI technology and modern software engineering.
- Introduce students to an increasingly important research area, thus motivating them to pursue more advanced courses in machine learning and to pursue undergraduate research projects in this area.

These objectives are being accomplished through the development, implementation, and testing of a suite of adaptable and self-contained, hands-on open laboratory projects that can be closely integrated into the AI course. In addition to the lab manual, an instructor's manual that provides a sample syllabus to go with each project and guidance for the adoption and adaptation of this curricular material will also be compiled. The instructor manual will include sample solutions as well as supporting material such as code and documentation. All curricular material with the exception of solutions is being made available at the project website at <http://www.removedforanonymity>.

BRIEF OVERVIEW OF PROJECTS

We have identified the following applications and learning models for our projects. This section provides a brief overview of each of the projects. Detailed descriptions along with contact information of the developer of each project are available at the project website.

Web User Profiling

Web searches provide large amounts of information about web users. Data mining techniques can be used to analyze this information and create web user profiles. A key application of this approach is in marketing and offering personalized services, an area referred to as "data gold rush". The aim of this project is to develop a system that can be used to develop an intelligent web browser. This project focuses on the use of Decision Tree learning to create models of web users.

Character Recognition and Learning with Neural Networks

The power and usefulness of artificial neural networks have been demonstrated in several applications including speech synthesis, diagnostic problems, medicine, business and finance, robotic control, signal processing, computer vision

and many other problems that fall under the category of pattern recognition. The goal of this project is to develop a character recognition system based on a neural network model.

Solving the N-Puzzle Problem

The N-puzzle game provides a good framework for illustrating conceptual AI search in an interesting and motivating way. The objective of this project is to introduce the student to Analytical (Explanation-Based) Learning using the classical AI framework of search. Hands-on experiments with search algorithms combined with an Explanation Based Learning (EBL) component give students a deep, experiential understanding of the basics of EBL.

Solving the Dice Game Pig

The jeopardy dice game Pig is very simple to describe, yet the optimal policy for play is far from trivial and was only recently solved. Using the computation of the optimal solution as a central challenge problem, we give the student a deep, experiential understanding of dynamic programming and value iteration through explanation, implementation examples, and implementation exercises.

Web Document Classification

Along with search engines, topic directories are the most popular sites on the Web. Topic directories organize web pages in a hierarchical structure according to their content. The aim of the project is to investigate the process of tagging web pages using the topic directory structures and apply Machine Learning techniques for automatic tagging. This would help in filtering out the responses of a search engine or ranking them according to their relevance to a topic.

The Game of Clue

The popular board game Clue serves as a fun focus problem for this introduction to propositional knowledge representation and reasoning. After covering fundamentals of propositional logic, students first solve basic logic problems with and without the aid of a satisfiability solver. Students then represent the basic knowledge of Clue in order to solve a Clue mystery. Could the current best stochastic SAT solver, Adaptive Novelty+, benefit from machine learning elements?

Support Vector Machine to Analyze DNA Microarrays

A support vector machine (SVM) is a powerful machine learning technique that is used in a variety of data mining applications, including the analysis of DNA microarrays. A DNA microarray is a small silicon chip that is covered with thousands of spots of DNA of known sequence. Biologists use microarrays to study gene expression patterns in a wide range of medical and scientific applications. The goal of this project is to learn how to use an SVM to recognize patterns in microarray data. Using publicly available data sets, we will train an SVM to distinguish between two forms of leukemia.

Genetic Algorithms

Anyone familiar with the theory of natural selection can imagine how a difficult search problem might be attacked by combining the elements of reasonably good solutions and randomly mutating the resulting "offspring" to preserve variety. At a deeper level, there is often a disconnect between real-world problems and the kind of problems that students work on in an undergraduate AI or machine-learning course. This project focuses on learning and applying two recently developed genetic algorithms for solving real-world problems in multi-objective optimization and evolving behaviors for competitive robotic agents.

Biomedical Term Classification

Due to the explosive growth of knowledge in biotechnologies (about 1500 research abstracts are added every single day to MEDLINE, an electronic repository of biomedical papers) an acute need for knowledge management tools has arisen. Natural Language Processing (NLP) can help fulfilling this acute need. A major problem in BioNLP, the area of research at the intersection of Biotechnologies and NLP, is that same biomedical term can be frequently used with different meanings in biological texts. For instance, SBP2 can refer both to a protein or a gene. In this project, we combine NLP with Machine Learning techniques, namely decision trees and Naïve Bayes, to build software tools that classify biomedical terms based on the surrounding contexts in which they appear. In particular, we work with terms that refer to the following categories: DNA, RNA, protein and cell_line, cell_type.

General-Purpose Problem Solver

Genetic programming (GP) is perhaps the most general of local search algorithms. It is particularly useful in solving design and optimization problems. Strengths of GP include the ability to work with heterogeneous data, and that a relatively low amount of information needs to be specified to achieve success. A number of patents now exist that have been achieved using genetic programming. The goal of this project is to learn about machine learning (problem formulation, search, and knowledge representation) by building a basic genetic programming framework and to use it to solve problems. The framework will be built piece-by-piece, as concepts are introduced.

Supervised Learning of Sign Language Characters

Recognition of images is a key technological advancement that has come out of the field of artificial intelligence. Machine learning technologies can be used to learn and recognize a variety of objects contained in images. The results can be used for face recognition, character recognition, gesture recognition, and a range of additional applications. The goal of this project is to train several different classification algorithms to recognize the alphabet character that is being signed using the American Sign Language (ASL) gesture.

Probabilistic Reasoning with Naïve Bayes and Bayesian Networks

Bayesian (also called Belief) Networks (BN) are a powerful knowledge representation and reasoning mechanism. BN represent events and causal relationships between them as conditional probabilities involving random variables. Given the values of a subset of these variables (evidence variables) BN can compute the probabilities of another subset of variables (query variables). BN can be created automatically (learned) by using statistical data (examples). The well-known Machine Learning algorithm, Naïve Bayes is actually a special case of a Bayesian Network.

The project allows students to experiment with and use the Naïve Bayes algorithm and Bayesian Networks to solve practical problems. This includes collecting data from real domains (e.g. web pages), converting these data into proper format so that conditional probabilities can be computed, and using Bayesian Networks and the Naïve Bayes algorithm for computing probabilities and solving classification tasks.

Relational Learning for Web Document Classification

Most of the content-based approaches to text and web document classification are based on the bag of words model, well known from the area of Information Retrieval. This model is simple and efficient, but fails to capture many additional document features such as the internal HTML structure, language structure and inter-document link structure. All this however may be a valuable source of information for the classification task. The basic problem with incorporating this information into the classification algorithm is the need for uniform representation. For example, the content-based classification works well with the vector space representation, while hyperlink-based classification can be implemented by using graph models. This project introduces students to Relational (First-Order) Learning that allows various kinds of information to be represented in a uniform way and used for document classification. One of the most successful relational learning systems, FOIL is used to create relational representation of web documents and to solve classification problems.

Machine Learning for Automated Reasoning

Automated theorem proving (ATP) is concerned with the development and use of systems that automate sound reasoning: the derivation of conclusions that follow inevitably from facts. A key concern of ATP research is the development of more powerful systems, capable of proving more difficult theorems. Automated reasoning in large theories is becoming more prevalent as large knowledge bases are translated into forms suitable for ATP. Of particular interest is to improve performance of ATP when solving many problems from the same axiom set, but each problem requires use of only a subset of the axioms. The focus of this project is the extraction of a sufficient subset of axioms for proving that a given conjecture is a theorem, using machine learning (ML) to learn from existing proofs which axioms are more likely to be used in a proof of the theorem.

Robot Defense: Intelligent Behavior in a Real-Time Strategy Game

Interactive real-time games form the basis of many modern research and simulation environments. These dynamic environments allow a wide range of real-world issues to be explored such as the how computation time can negatively impact decision making. In this project, students provide the AI for an interactive real-time game called Robot Defense (similar to the highly popular Flash game Desktop Tower Defense). Assignments leverage search, knowledge representation and reinforcement learning to create low level controls for game entities and high level decision making for an agent that can play the game autonomously.

Generating Dense Boggle Boards with Genetic Algorithms

The problem of generating dense boards (high-scoring boards) in the game of Boggle is a challenging and fun task. This project explores the fundamentals of genetic algorithms, which are well suited for this task for two reasons: first, it is impractical to exhaustively examine or maintain all candidate boards; and second, information gained along one local search path can be integrated in multiple meaningful ways with information gained along other paths.

Application of Associative Matrices to Recognize DNA Sequences in Bioinformatics

Associative matrices are considered a type of neural network topology used to recall and recognize previously known or unknown patterns. For example, the Hebbian linear associative matrices can be trained to recognize a particular DNA sequence into another specimen sequence that may help biologist to identify similarities and other characteristics important in the knowledge and recognition of a particular sequence in another specimen. This approach may result especially beneficial in mutated sequences where mutations or other changes in the sequence as deletions and insertions are present. Associative matrices have been used to recognize characters, shapes, or specific objects from an image. Such changes are considered noisy patterns that are one of the important features of using associative matrices in this field.

The project introduces students to the use of associative matrices concepts in learning and pattern recognition. Students will experiment the use of the matrices to recognize DNA sequences and its possible mutations.

Creating Robotic Intelligent Agents

An intelligent agent approach to the understanding of basic AI principles has become widely accepted. This project uses robotics to create simple reflex, model based, goal based, utility, and learning agents. The simple reflex, model based, goal based, and utility robot agents are programmed to play a modified soccer game. The learning robot agent uses a neural network to learn line following behavior.

Competitive Learning in Checkers

One of the earliest investigations of machine learning was Arthur Samuel's work with Checkers, developing a program that learned a board evaluation function through the experience of playing many games. Starting from Samuel's

approach, students will develop Checkers programs that learn to play by competing with other students' programs. Competitions will take place over the internet using a simple server that communicates moves between clients, checks for the validity of moves, keeps track of (and limits) the time spent by each player, and reports the winners and losers.

In addition to the reinforcement learning aspects, the project deals with the problems of search, efficient representation of knowledge, and reasoning about time as a resource. Using the same server guarantees that each checkers program learns and can compete within the same environment. At the end of the semester, the project teams compete in a tournament, which adds yet another level of excitement and provides a straightforward, if relative, comparison of the resulting programs.

Route Planning, State-Space Search, and Case-Based Learning

Route planning is the process of determining a route, or path, to follow to move from a starting location to a goal location. It is a good application to be tackled by simple state-space search methods, at least on the small scale. Re-generating routes from scratch can be expensive, and seems wasteful. Therefore case-based reasoning will be explored as a method for storing and re-using previously generated routes.

Case-based reasoning systems store information, route plans, in this case, according to a set of features that describe the context in which they apply. For route plans, the context might be just the starting and goal locations, but could also include more elaborate features like passing particular obstacles, going past particular landmarks, or being scenic or fastest. This project introduces students to the costs and benefits of state-space search algorithms, and examines CBR as an alternative or adjunct to a brute-force method. It experiments with how to break route plans into cases, how to choose useful indices, how to judge similarity, how to adapt a partial-match to fit a new context, and how to learn by storing new cases in the CBR memory.

Solving the Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is a classic optimization problem: find the least-cost round-trip route amongst N cities, visiting all cities exactly once. In this set of assignments students will develop alternative machine learning approaches to solving the TSP. Some of the possible techniques include: genetic algorithms, simulated annealing, Boltzmann machines, Kohonen SOM, ant colony systems etc. The learning goal of these assignments is to use the TSP as a test-bed problem to study various machine learning techniques individually and comparatively.

Recognizing Images of American Sign Language Letters Using Principal Component Analysis

Digital images are now ubiquitous and easy to acquire. While humans easily recognize the objects and other semantic content in images, it has been much more difficult to do so automatically. To obtain higher-level semantics, a set of features must be computed from each image for comparison to the features of the reference image. Finding the right set of

features by trial and error can be time-consuming and difficult. Therefore, in this project, we will use Principal Component Analysis (PCA) to learn an appropriate set of features from a training set of images, and then apply them to a test set of images. The methodology used in this project has been used successfully for face recognition.

Discovering Optimal Policies with Value Iteration, Policy Iteration, and Q-learning

Value and Policy Iteration provide an excellent means for agents in a nondeterministic environment to determine an optimal series of actions through the solving of a Markov decision process (MDP). However, solving an MDP requires that an agent have a great deal of knowledge about its environment: specifically, the rewards for each state and the transition probabilities between states. When this knowledge is not available to the agent, it can be learned through experience. Reinforcement learning, specifically Q-learning, is a method for doing this. Q-learning is a form of model-free learning; a Q-learning agent can learn an optimal policy without any knowledge about its environment, given enough experience. In this problem, students implement value iteration, policy iteration, and Q-learning to discover optimal policies for both a toy map and for a realistic campus map. Applying these two approaches to the same set of problems provides students with an understanding of how learning can be used to make up for a lack of available domain knowledge.

Understanding Genetic Algorithms through Interactive Play

Genetic algorithms explore the hypothesis space in search of a best hypothesis to solve a problem through the manipulation of a string representation of those hypotheses. Inspired by biological evolution, parts of good hypotheses are combined and used with limited exploration to evolve new hypotheses similar to how different genes could potentially be combined into better chromosomes making a better creature. This project presents an interactive game environment where the player (student) has to evolve members of a tribe of creatures, Modabus, in order to have the correct genetic makeup to overcome an obstacle impacting their society. Through manipulation of crossover, mutation, and selection the player (student) will guide the genetic algorithm to find a solution (a Modabu citizen with the proper level of intelligence, strength, and speed to overcome each obstacle) over the course of three increasingly difficult levels.

DISCUSSION AND EXPERIENCES

The effectiveness of this project in achieving each of the goals listed earlier is being evaluated through a multi-tier evaluation system involving (1) the students taking the AI course, (2) members of the advisory board, and (3) the two PIs and 18 participating faculty in the development and testing of the projects. Led by the project evaluator, the evaluation process involves both a formative evaluation and a summative evaluation. The formative evaluation includes strategies for

CONCLUSION

Question	Average	A or SA
6) Project contributed to my overall understanding of material.	3.09-5.0	75-100% []
12) After taking this course I feel that I have a good understanding of the fundamental concepts in AI.	4.0-4.69	75-100% []
16) Based on my experience with this course, I would like to learn more about the field of AI.	4.18-5.0	73-100% []
22) I had a positive learning experience in this course.	4.18-5.0	75-100% []

Figure 1: Evaluation Results

monitoring the project as it evolves and providing feedback to guide the development efforts. An assessment of the project's effectiveness in improving student learning and an evaluation of our efforts at different stages of the project will be used to assess progress made toward our goals and will direct the development of our work. The summative evaluation involves strategies to evaluate the effectiveness of the curricular material and our work in achieving our goals and in identifying findings at the end of the project.

We have completed the development of draft copies of twenty-six projects, which are now available at the project website. During the last year, we started implementing the projects in courses at the institutions of the faculty developers. We will continue the implementation during spring 2009 and in the upcoming semesters. Feedback from the earlier implementations is being used to further refine the projects.

Students were surveyed using a student survey comprised of 22 ranked questions and two open-ended questions. In addition, a sample of students was interviewed from each institution using either a direct (group) interview or a group conference call. On the written student evaluation, items were scored from 1 to 5 with 5 being the highest level. Preliminary surveys of students as well as student interviews revealed a high level of satisfaction with the course. Figure 1 presents results of four questions of the student survey administered and a summary of student responses, including the averages and range of percentages of those who replies with 'agree' or 'strongly agree'. These represent results of thirteen course offerings that used the projects during the last year. The number of students per class varied from 10 to 25 students.

The general feeling about using the projects in classes is positive. In these initial experiences, students revealed a high satisfaction with the course. They liked being able to apply the problem-solving techniques to a "real" situation and to see how they worked. They also liked the opportunity to choose their own particular application domain and further explore or extend the projects in areas they found interesting and challenging. While we have mostly used the projects in the introductory artificial intelligence courses, they have also been used in other upper level courses such as software engineering and capstone courses. Following the course, several students enrolled in independent study course to extend the work of the project.

To discuss our experiences and to support others in adapting our projects, we are in the process of planning for an NSF-funded workshop as a pre-conference activity. We received significant interest in the workshop with 45 applications submitted for only 20 slots.

Further testing of the material is currently underway.

We presented curricular material that incorporates machine learning as a unifying theme to teach fundamental concepts typically covered in the introductory artificial intelligence courses. We presented preliminary experiences. The projects were well received by the students. By using projects involving real-world applications, we provided additional motivation for students. While illustrating core concepts, the projects introduced students to an important area in computer science, machine learning, thus motivating further study. Feedback from these earlier implementations is being used to further refine the projects. Additional implementation and testing is currently underway that will further investigate the impact of theme and problem based learning on student outcomes and examine whether the results of earlier phases can be extended and generalized by other faculty across the country.

ACKNOWLEDGEMENT

This work is supported in part by National Science Foundation grant *DUE-0409497* and *DUE-.0716338*

REFERENCES

- [1] Dodds, Z. et al. (eds.), Robots and Robotics in Undergraduate AI Education, Special Issue of AI Magazine, 27(1), AAAI Press, 2006.
- [2] Fox, S. Introductory AI for Both Computer Science and Neuroscience Students, Proceedings of the 20th International FLAIRS Conference, AAAI Press, May 2007.
- [3] Greenwald, L. (ed.), Accessible Hands-on Artificial Intelligence and Robotics Education, AAAI Press Technical Report, March 2004.
- [4] Harlan, R., Levine, D., and McClarigan, S., The Khepera Robot and the kRobot Class, Proceedings of the thirty-second SIGCSE technical symposium on Computer Science Education, ACM Press, New York, NY, February 2001, 105-109. DOI=<http://doi.acm.org/10.1145/364447.364553>
- [5] Hearst, M. (ed), Improving Instruction of Introductory AI, Tech. Report FS-94-05, AAAI Press, 1995.
- [6] Kumar, A., et al., Non-Traditional Projects in the Undergraduate AI Course, Proceedings of the 37th Annual SIGCSE Technical Symposium on Computer Science Education, ACM Press, March 2006. DOI=<http://doi.acm.org/10.1145/1121341.1121491>.
- [7] Removed for anonymity, "Thrown Down an AI Challenge", Proceedings of AAAI Spring Symposium on Using AI to Motivate Greater Participation in Computer Science, AAAI Press, March 2008.
- [8] Removed, for Anonymity, Pedagogical Possibilities for the N-Puzzle Problem, Proceedings of the 36th Frontiers in Education Conference, IEEE Press, November 2006.
- [9] Mitchell, T., The Discipline of Machine Learning, CMU-ML-06-108, July 2006.
- [10] Nilsson, N., Artificial Intelligence: A new Synthesis, Morgan Kaufmann, Publishers, 1998.
- [11] Removed, for Anonymity, Unifying Undergraduate AI Courses through Machine Learning Projects, Proceedings of the 25th ASEE Conference, June 2005. DOI=
- [12] Russell, S. and P. Norvig, Artificial Intelligence: A Modern Approach, Author's Website, <http://aima.cs.berkeley.edu/>.