

A Survey of Current Practice and Teaching of AI

Michael Wollowski,¹ Robert Selkowitz,² Laura E. Brown,³ Ashok Goel,⁴
George Luger,⁵ Jim Marshall,⁶ Andrew Neel,⁷ Todd Neller,⁸ Peter Norvig⁹

1) Rose-Hulman Institute of Technology, wollowski@rose-hulman.edu, 2) Canisius College, selkowitz@canisius.edu,
3) Michigan Technological Institute, lebrown@mtu.edu, 4) Georgia Institute of Technology, ashok.goel@cc.gatech.edu,
5) University of New Mexico, luger@cs.unm.edu, 6) Sarah Lawrence College, jmarshall@sarahlawrence.edu,
7) Discover Cards, ajneel@acm.org, 8) Gettysburg College, tneller@gettysburg.edu, 9) Google, pnorvig@google.com

Abstract

The field of AI has changed significantly in the past couple of years and will likely continue to do so. Driven by a desire to expose our students to relevant and modern materials, we conducted two surveys, one of AI instructors and one of AI practitioners. The surveys were aimed at gathering information about the current state of the art of introducing AI as well as gathering input from practitioners in the field on techniques used in practice. In this paper, we present and briefly discuss the responses to those two surveys.

Introduction

During the 2014 EAAI conference, Todd Neller conducted an informal survey of the participants about which topics they teach in their AI courses as well as which topics they would like to teach. The survey suggested an apparent disparity between what is taught in many AI courses and what many AI colleagues would like to teach. We decided to study this issue by conducting surveys of both educators and practitioners.

There have been at least two events in the past twenty or so years which looked at the challenges of AI Education: the 2008 “*AI Education Colloquium*” held at the 2008 AAAI conference (AAAI 2008) and the 1994 AAAI Fall Symposium entitled “*Improving Instruction of Introductory Artificial Intelligence*.” (AAAI 1994) This is in addition to regularly held *Educational Advances in AI Symposia* collocated with the AAAI conference and past *AI Education* tracks at FLAIRS. To quote Marti Hearst, the program chair of the 1994 AAAI Fall Symposium (Hearst 1994), “This symposium was motivated by the desire to address the oft-voiced complaint that introductory artificial intelligence is a notoriously difficult course to teach well.” With the regular progression of the field and recent successes such as autonomous cars and IBM’s Watson cognitive system, this situation has not become easier. We believe that

now is a good time to reevaluate the practice of teaching Artificial Intelligence.

In this paper, we present and evaluate two surveys. One was aimed at practitioners and the other one at educators. Broadly, the practitioner’s survey was aimed at gathering information about which topics and techniques are being used in practice. Additionally, we invited practitioners to offer their insights into what topics they believe should be taught as well as what sort of course outcomes would be desirable. This sort of data is routinely sought by accredited institutions aiming at obtaining feedback from their stakeholders.

The aim of the educator’s survey was to capture current practice of teaching AI as well as gather input on proposed changes. We gathered information about the kinds of institutions at which our colleagues teach, prerequisites and course outcomes as well as detailed information about which topics are being covered, which topics colleagues would like to cover, and reasons for any disparities between these topic lists.

Surveys

In order to establish current practice and desires, we conducted a survey of AI educators and a separate survey of AI practitioners.

Educator’s Survey

We invited AI educators to participate in an online survey. In order to obtain precise input, we requested data on the kind of institution our colleagues teach at, prerequisites and learning outcomes for their AI course as well as major topics covered. Furthermore, we gathered input on desired topics and reasons for non-inclusion as well as what topics should be taught.

In order to manage the input, we asked our participants to list topics under thirteen categories. The categories given are those that are used on the web site aitopics.org maintained by AAAI.

1. Please describe the kind of institution at which you teach. Select all that apply.
 - Undergraduate only in CS
 - Masters program in CS
 - Doctoral program in CS
 - Large (above 3000 students)
 - Small
 - Liberal arts
 - Science or engineering
2. Please list the prerequisite topics for your AI class.
3. Please list the primary goals/outcomes for your AI course.
4. Please indicate any other AI related courses offered at your institution. Select all that apply.
 - None
 - Graduate AI course
 - Machine Learning
 - Robotics
 - Natural Language Processing
 - Other (please specify)
5. Please enter the major topics covered in your class. List items in the appropriate text-boxes.
 - Applications
 - Cognitive Science
 - Ethics & Social Issues
 - Games & Puzzles
 - History
 - Machine Learning
 - Natural Language
 - Philosophy
 - Representation and Reasoning
 - Robots
 - Speech
 - Vision
 - Others
6. Please enter major topics that you would like to cover but currently do not. List items in the appropriate text-box.

[Same list of categories as for question 5]
7. What is the major impediment to covering those topics?
8. More broadly, which AI topics/techniques do you think should be covered in an introductory AI course? Please list them in order of importance.
9. Out of curiosity, what is your definition of AI?
10. Any last comments?

The educator's survey was distributed through the mailing lists of AAI, SIGCSE, SIGAI, the EAAI-13 participants, the organizing committee of the FLAIRS AI education track, and ml-news. It was also made available on the EAAI-16 website and sent to contacts in academia.

Practitioner's Survey

We were interested in learning about modern, relevant, and cutting-edge AI topics for potential inclusion in an introductory AI course. We invited AI practitioners from industry, government and academia to participate in an online survey. Below are the questions we asked.

- 1) Which AI topics/techniques do you use in your work?
- 2) Which AI topics/techniques do you consider up and coming?
- 3) Which AI topics/techniques do you think should be covered in an introductory AI course? Please list them in order of importance.
- 4) More broadly, which learning outcomes do you think an introductory AI course ought to have?
- 5) Out of curiosity, what is your definition of AI?

The practitioner's survey was distributed through the mailing lists of AAI, EAAI-13 participants, the organizing committee of the FLAIRS AI education track and ml-news. It was also made available on the EAAI-16 website and sent to contacts in industry.

Evaluation of Educator's Survey

We received 59 responses to this survey, though only 37 respondents went beyond the first four questions. Since the most valuable data was collected through questions five to eight, we decided to evaluate only those 37 responses that had some data beyond question 4. Due to space limitations, we will provide an analysis of only the most important results. A more fine-grained analysis will be provided in another venue.

Institutional Background

Through Question 1, we attempted to obtain information on the background of our respondents. Since multiple submissions were allowed, we decided to only count the highest degree listed. Based on this criterion, about half the respondents (43%) teach at institutions that offer a bachelor's degree as the highest CS degree. About a fifth (19%) offer a master's degree and about a third (30%) are from CS Ph.D. granting institutions. As such, we have a balanced mix of participants from across the spectrum of primarily teaching to primarily research institutions.

When it comes to the size of the institutions, we are just about evenly split between those with a student body above 3000 students (41%) and those below (35%). 24% of the respondents did not answer this question.

A similar split exists between liberal arts colleges (41%) and science or engineering schools (35%). It should be noted that not everyone answered this question and some respondents selected both options.

Prerequisites to Existing AI Courses

In Question 2, we gathered input on the prerequisites for AI courses. Many require data structures as a prerequisite (46%). Equally many may assume a year-long software development sequence. Some responses contained the term “algorithms” (19%). It was not clear whether this prerequisite is concerned with the study of complexity or with a more practical introduction containing a fair amount of software development. The wordings of some of the responses suggest it is the latter.

A good number of our colleagues may assume a mathematical background in the form of a discrete mathematics course (27%) or a course in probability or statistics (19%).

On the other side of the spectrum, 11% of AI courses do not have prerequisites.

Primary Goals and Outcomes

This third question was difficult to evaluate. We were not able to properly code seven of the responses. The evaluation in this section ignores those seven responses. See Figure 1 for a graphical representation of the data.

As expected, an introductory course should cover the basic or main topics of artificial intelligence (77%). Many respondents used those keywords outright; others provided lists of topics that typically appear in the first couple of chapters of several AI textbooks (Russell 2010, Luger 2009).




Figure 1: Primary goals and outcomes

Thirteen percent of the respondents provide a survey of AI techniques. Based on the wordings of the responses, the line between the “basics” and “survey” categories is rather blurry. If we add the values of those two categories, then 90% of respondents cover the basic topics.

45% of the courses either explicitly mention a software development component or use language that is strongly suggestive of such a component.

16% of courses cover topics about the societal impact, ethical concerns or philosophical issues of AI.

Other AI Courses

Question 4 inquired about additional AI related courses offered at our colleague’s institutions. 59% of the survey participants offer a machine learning course, 46% offer a robotics course and 35% offer a course on natural language processing.

38% of institutions offer graduate courses. This makes sense, given that 53% of respondents come from institutions that offer a masters or Ph.D. in CS.

46% of institutions offer some other AI related course. Only 14% of institutions have no other AI courses. All except for one who do not offer additional AI related courses are undergraduate only institutions (4). However, many undergraduate only institutions (12/16) offer a wide selection of AI related courses.

Topic Coverage in AI Courses

Questions 5-7 of the survey addressed issues of topic coverage in introduction to AI courses: which topics are covered (Q5), which topics instructors wish to cover but cannot (Q6), and the reasons why instructors do not cover all desired topics (Q7).

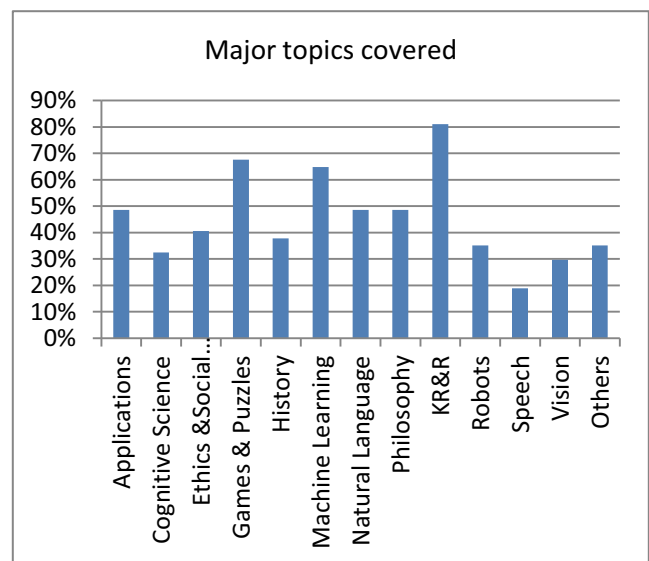


Figure 2: Major topics covered in AI courses

Consistent with course outcomes, knowledge representation and reasoning (KR&R) gets top billing with 81%. This is not a surprise as this category, at least according to ai-topics.org (AITopics 2015) includes search. Machine learning, and games and puzzles each were mentioned by 68% of the instructors. Judged by the responses to Question 8, where games were mentioned by just 11% of the respondents, puzzles and games are generally not consid-

ered as core topic of an AI course; however they undoubtedly keep our students' interest. Another reason for the high frequency count of this topic is perhaps that they are a fine way to introduce a variety of search techniques and are likely used in that way. A similar argument can be made for the high coverage of applications (49%); they too are a fine way to motivate students and they serve to satisfy learning outcomes related to strengthening students' software development skills.

41% of our colleagues cover ethics and social issues and 49% cover philosophical issues.

As can be seen in Figure 2, the range of topics instructors are offering is quite large. Clearly any attempt to cover all the listed topics in one term would not do justice to them. As such, we see that AI courses vary widely in both the number and selection of topics to be covered.

Desired Topics

Perhaps the most interesting result of this section is that about a third of the instructors (35%) do not wish to add topics to their AI course. While machine learning is cited most frequently, we again see a wide spectrum of topics that our colleagues wish to add. Looking at the actual responses, we do not see any strong trends. The responses seem to be based on the desires of individual instructors.




Figure 3: Desired topics

Impediments to Adding Topics

The single most pressing constraint which prevents the addition of topics to the AI curriculum is time (69%). The lack of suitable learning materials does not seem to be a factor as it is mentioned by only 15% of our respondents. Twelve percent of our colleagues cited lack of student preparation or interest as a reason for not adding desired topics.

Recommended Topics

While we presume that instructors cover the topics that they deem important, there may be constraints that prevent them from following their wishes. As such, we inquired about what topics and techniques ought to be covered. Lending insight into what are considered basic topics such as mentioned in Question 3; we see that an ideal course covers search, knowledge representation and reasoning as well as machine learning. It should be noted that this was an open-ended question; as such "search" was not combined with "knowledge representation and reasoning".

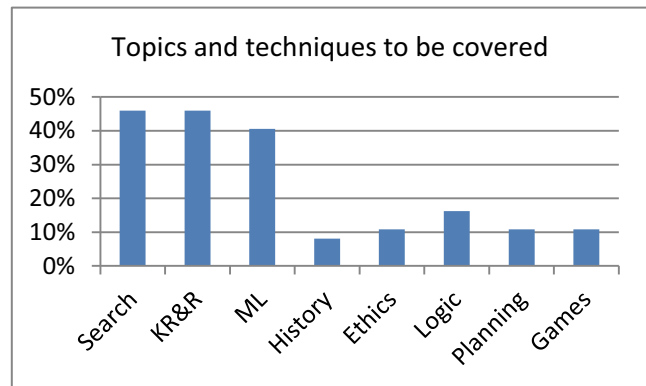


Figure 4: Topics and techniques that should be covered

Definition of AI

Primarily out of curiosity, we invited our respondents to tell us about their definition of AI. If we eliminate non-responses (7), then about half of the responses (43%) were along the lines of producing software that exhibits traits that we find in humans. About a tenth of the responses (11%) took a view of AI as producing goal oriented agents. Eight percent focused their definition of AI as solving the next big challenge. About a fifth (19%) gave various other definitions.

Comments

We invited instructors, to provide any further feedback they may have. Four respondents mentioned something to the effect of how difficult it is to teach an AI course, in particular with regard to choosing which of the many topics to cover. In addition, we received several encouraging emails about our project, indicating that collecting and evaluating the surveys is of interest to our community.

Evaluation of Practitioner's Survey

This survey was taken by 31 respondents.

Topics and Techniques in Use

We started out this survey by asking which topics and AI techniques the respondents are using in their work. We took the list of categories given to instructors and coded the open-ended responses of the practitioners.

Again knowledge representation and reasoning tops the list with 61%. Perhaps not surprisingly, applications are next with 55%. After all, this is what practitioners in industry do. Machine learning is a close third (52%) and natural language processing is mentioned by a quarter of the respondents.




Figure 5: Topics and techniques used by practitioners

Up-and-coming Topics and Techniques

We were curious to learn from practitioners about which topics they see as up-and-coming.

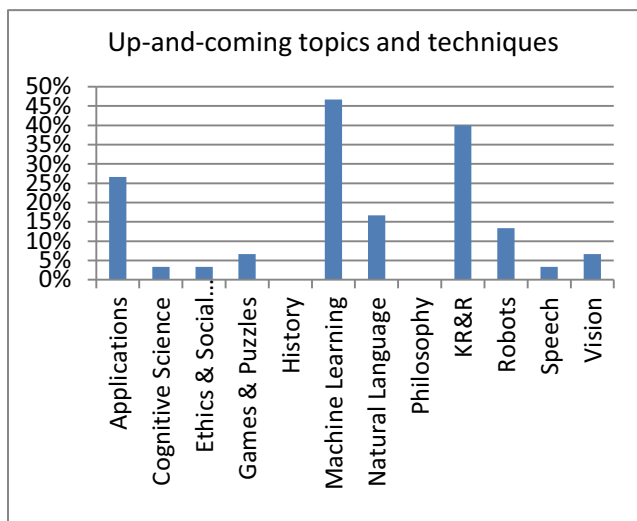


Figure 6: Up-and-coming topics and techniques

Machine learning tops the list with 47%, knowledge representation and reasoning is second with 39%, and applications are third (27%).

Two major areas of current work, natural language processing and robotics are also seen as important with 17% and 13% scores, respectively.

We point out that there is a role reversal of the top two items between what practitioners do on a daily basis and what they see as up-and-coming, namely between knowledge representation and reasoning, and machine learning.

We note that various applications are also seen as on the rise. Perhaps, this is an indication of a desired focus on engineered systems.

Recommended Topics and Techniques

Attempting to look ahead, we asked practitioners about which topics and techniques ought to be covered.

When analyzing the responses, we eliminated those that suggested an amount of topics that cannot reasonably be covered in an introductory course (5 responses) as well as responses that could not be matched to the AI topics list even under generous interpretations. The remaining responses were categorized to the AI topics list used in the instructor's survey.

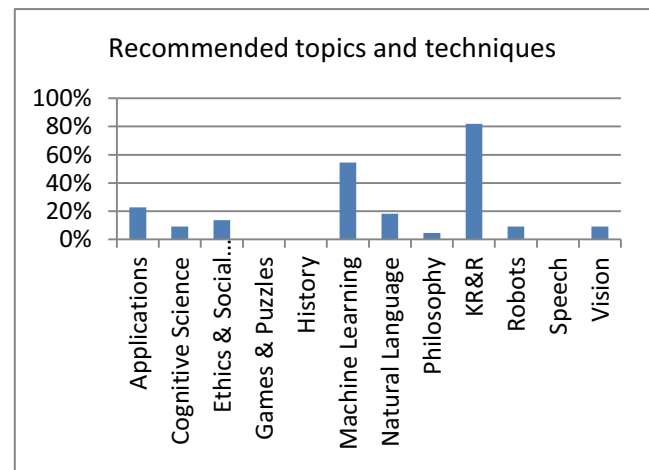


Figure 7: Topics and Techniques recommended by practitioners

Mirroring the data of the first question, knowledge representation and reasoning ranked first (82%) and machine learning second (55%). Applications and natural language processing are third and fourth with 23% and 18%. Ethics was seen as important by 12% of the respondents.

Recommended Learning Outcomes

When asked about learning outcomes, we purposefully provided the respondents with a blank canvas. We reaped the rewards of this choice by obtaining a wide and insight-

ful set of responses. We categorized the responses based on patterns that emerged.

Almost half of the respondents (48%) would like us to lay a good foundation or introduce basic tools and techniques. Almost the same number of respondents (44%) would like our students to have the ability to take perspective on AI and its methods so as to evaluate their strengths and weaknesses. We did not see this in the instructors' responses. We also did not ask for that information explicitly, and perhaps this is something that should be asked in a follow-up survey. Forty-one percent of the respondents would like our students to have skills that we interpreted as having the ability to engineer a system. While our students get exposed to this through software development exercises, systems engineering is likely not a major outcome of those assignments. About a quarter of the respondents (22%) would like our students to be exposed to cutting-edge work. This is in contrast to the 11% of the courses that have this item as an outcome.

The picture that we see as emerging here is that practitioners need to have basic skills as well as the skill to engineer a system and the ability to critically evaluate tools and techniques, whether existing or at the cutting edge.

We saw some combinations of responses that are noteworthy. When "engineering a system" is mentioned, 7 out of 11 times, it is mentioned all by itself. Three out of 11 times, it is mentioned with just the ability to "evaluate" a system.

Definition of AI

Primarily out of curiosity, we invited our respondents to tell us about their definition of AI. The responses by the practitioners were very homogenous. Two did not provide a definition and two stated that they prefer not to provide one. All other responses could be neatly placed in four categories. About half of the responses (48%) were along the lines of producing software that exhibits traits that we find in humans. About a quarter of the responses (26%) took a view of AI as producing goal oriented agents; this is about twice the percentage of the instructors' survey. Another quarter (26%) focused their definition of AI as building sophisticated or complex systems. This is something that we did not see in the educators' survey and is perhaps suggestive of the fact that many practitioners are focused on building artifacts that perform. Only 4% of the submissions defined AI as doing what computers cannot do yet.

Conclusions

Through our surveys, we documented current practice and teaching of AI.

In summary, there seems to be a healthy match between instructed topics and the practitioners' needs. The big three:

search, knowledge representation and reasoning, and machine learning are well represented in the results of both surveys.

While there are many commonalities, some differences stand out. Practitioners, through their input on learning outcomes and their definition of AI, suggest systems engineering as a major learning outcome. This is something we did not see in the educators' data. Instead, the games and puzzles category looms large in the educators' survey responses. While for practitioners, it might be desirable that we expose our students to problems that require a strong systems engineering component, for educators "toy problems" such as puzzles and games are a good way to give a concise introduction to basic AI topics, in particular to students who by and large have only a year or two of software development expertise.

Another discrepancy concerns the ability to take perspective of AI tools and techniques. This is an important skill to practitioners; however, it is not something that educators seem to cover in detail. Instead, instructors seem to take a broader perspective, focusing on historical, ethical and philosophical issues instead.

Acknowledgements

We thank AAI, in particular Carol Hamilton, for helping us conduct the two surveys.

References

- AAAI 1994. Improving Instruction of Introductory Artificial Intelligence, in: *Proceedings of the 1994 AAI Fall Symposium*, Menlo Park, CA.: AAI Press.
- AAAI 2008. AI Education Colloquium, in *Proceedings of the 2008 AAI Conference*, Menlo Park, CA.: AAI Press.
- AITopics 2015. <http://www.aitopics.org>
- Hearst, M. 1994. Preface: Improving Instruction of Introductory AI. in: *Proceedings of the 1994 AAI Fall Symposium*, Menlo Park, CA.: AAI Press.
- Luger, G. 2009. *Artificial Intelligence, 6th Ed.*, Boston, MA.: Addison-Wesley.
- Russell, S., and Norvig, P. 2010. *Artificial Intelligence, 3rd Ed.*, Upper Saddle River, NJ.: Prentice Hall.