The Pedagogy of Artificial Intelligence: A Survey of Faculty who Teach Introductory AI

Harlan D. Harris*
Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL 61801

Sarah M. Kiefer
Department of Educational Psychology
University of Illinois at Urbana-Champaign
Champaign, IL 61820

Abstract

This paper present highlights of the results of a survey of instructors of the introductory Artificial Intelligence course at colleges and universities throughout the United States. Particularly notable results were seen in the Syllabus and Teaching Techniques sections of the survey, and those are reported here. Comparing responses from Doctoral-Extensive universities with other responses, the latter group sometimes tends to teach a syllabus that seems out-of-date relative to current AI research trends. Instructors from less research-oriented institutions reported spending significantly less class time on several current topics in AI, and significantly more time on several more outdated topics. Participants were also asked about what sorts of classroom teaching techniques they used, and what techniques were found to particularly valuable. In general, participants endorsed in-class activities to supplement lectures, and the use of programming homework, group projects, and an agent-themed syllabus.

Introduction

Introduction to Artificial Intelligence (ItAI) is a common computer science course offered by most 4-year colleges and universities offering bachelors degrees in Computer Science. As described by the ACM/IEEE Computer Science Curriculum (CS260) (Engel & Roberts 2001), ItAI is an upper-level undergraduate class, a survey of the field of Artificial Intelligence. The class has a reputation as being somewhat difficult to teach, with many different topics and a tendency to be rather disjoint. In addition to the basic theory and techniques used in AI research, instructors commonly cover the history and philosophy of AI, introduce specialized programming languages, and discuss various applications. This paper reports a subset of the results from a recently conducted survey of U.S. instructors of the ItAI class, focusing on issues related the syllabus and to effective teaching of the class. Complete results of the survey may be found in a technical report (Harris & Kiefer 2002).

In 1994, a AAAI workshop was organized on the topic of instruction of ItAI (Hearst 1994). With around three dozen faculty members in attendance, the workshop made several recommendations to instructors. In addition to advising that instructors organize the syllabus around a single theme (such as intelligent agents (Russell & Norvig 1994)), integrate AI with other areas of Computer Science, and emphasize programming and projects, the participants noted the significance of distinguishing "the old from the important". This conclusion, that some applications and techniques that were historically important are no longer worth teaching to introductory students, supports one of two main conclusions we will draw from the results of this survey.

The survey was designed to evaluate the current state of instruction of the ItAI course, and to answer questions raised by its reputation and by the AAAI workshop. Questions were asked about the students taking the course, the syllabus, the use of homework and the textbooks used, teaching techniques, organization and motivation of the course, and the faculty teaching the course. In this paper, we focus on the two sections with the most specific and practical results, the syllabus and teaching techniques.

The Methodology section reviews the data collection and analysis methods used, the Responses section presents selected responses, and the Conclusions section makes initial recommendations to faculty based on the data we received.

Methodology

The target population for the survey was people who have been primary instructors (not teaching assistants) of the introductory Artificial Intelligence class at colleges or universities in the United States, currently or in the past three years. We viewed departmental web pages for four-year institutions listed in the Carnegie Classification of Institutions of Higher Education (McCormick 2001). For each department, email addresses of faculty members who we thought may potentially teach the class were collected. (In some cases, this information was listed explicitly, but in general only faculty research interests, not teaching assignments, are available.) Note that only 10% of Masters, Liberal-Arts, and other Baccalaureate colleges were sampled, due to the large number of such colleges and the relatively low rate at which they offer the course. Consequences of this sampling are discussed below.

The survey was made available on the web for four weeks, and invitations to participate were emailed twice, first at the

*Correspondence concerning this paper should be addressed to Harlan D. Harris, who is now at Department of Psychology, Columbia University, 1190 Amsterdam Ave., MC 5501, New York, NY 10027. Email: harlan@psych.columbia.edu

beginning of the four week window, and again after two weeks. Efforts were taken to thank participants for their time and to ask that the invitation be forwarded to other members of the population whom we may have missed.

The survey itself was presented on the web. There were six sections, each a single web page, plus an introductory page at the beginning and a debriefing/thank-you page at the end. Most participants completed the survey in about 15 minutes. The survey may be viewed (data is no longer being collected) at: `http://langprod.cogsci.uiuc.edu/cgi-bin/itai-survey.cgi`. To the extent possible, we followed general guidelines for survey construction, including available guidelines for design of on-line surveys (Dillman, Tortora, & Bowker 1998). The survey was anonymous and voluntary, and participants were told that they could stop at any time during the survey or omit questions they preferred not to answer. The University of Illinois Institutional Review Board approved the survey.

Respondents

Table 1 shows the number and rates of responses. Overall, 404 institutions were sampled, and email addresses were found at 215 of those. We sent invitations to 479 valid email addresses, and received 102 responses. It is likely that a significant fraction of the population at the institutions surveyed was missed (undersampled), and it is also likely that a substantial fraction of those who received invitations were not actually members of our target population (oversampled). In addition, we have no way of knowing how many people who were not directly invited by us to participate may have received copies of the invitation, nor do we know the number of different institutions represented by the data collected. Due to these factors, any estimates of response rate must be viewed as only approximate. The 21% overall response rate that can be calculated from the numbers of invitations sent and responses received, however, is commonly viewed by survey researchers as a reasonably good rate for on-line surveys of this type.

More than 60% of the responses we received were from Doctoral-Extensive (formerly known as Research I) institutions. For many of the survey items, we broke down the results by comparing faculty who said they were at Doctoral-Extensive institutions ("Research" faculty) with those who said they were at other institutions ("Other" faculty).

Weighting and Statistics

As noted above, Masters-level and Baccalaureate-level colleges were sampled at 10%. To re-weight the data, the 15 responses from those institutions were weighted by a factor of 10, with the sample in question scaled so that the weighted *n* was not changed. When reweighting the entire set of data, responses from Masters and Baccalaureate institutions were weighted by 4.3, while other responses were weighted by 0.43. When comparing Research to Other responses, Masters and Baccalaureate responses were weighted by 2.2, Doctoral-Intensive, Technical, and responses that did not specify their institution type were weighted by .22, and Doctoral-Extensive institutions were not weighted. The data reported below is generally only weighted if the results were qualitatively different by doing so. Since all data analyzed statistically was categorical, the standard Chi-square test for significance was used, with $\alpha = 0.05$.

Responses

As noted above, only the results for the Syllabus and Teaching Techniques sections are presented here. Interested readers are referred to (Harris & Kiefer 2002) for details of the remainder of the results.

Syllabus

We asked respondents to, for a list of AI topics, specify whether they spend more than an hour, less than an hour, or no class time on each topic when teaching ItAI. The results shed light on questions regarding what topics are commonly taught, the amount that they are taught, and the extent to which the ItAI syllabus is consistent between universities.

We analyzed the results in two ways. First, we assigned each response a value of 1, 2, or 3, for Do not Discuss, $<$ 1 Hour, and $>$ 1 Hour, respectively, then sorted the topics by the mean results. The left-hand portion of Table 2 gives these results, which identify the most-common and least-common topics in the ItAI classroom. Second, we examined each topic, comparing Research respondents to Other respondents. The right-hand portion of Table 2 shows the percentages of responses in each of the three categories, either lumped together if there was no significant difference by Research vs. Other, or shown separately if there was a significant difference.

The overall ordering gives a general impression of which syllabus topics our respondents felt were most important. The variance is a rough measure of the consistency with which each topic is taught (or not taught). For example, Heuristics, with a high mean and low variance, are consistently taught extensively, while Semantic Networks, with a moderate mean and high variance, vary extensively to the extent to which they are covered. The extent to which variance in the syllabus may be viewed as a problem depends to a large extent on whether a particular topic is viewed as an important foundation for later coursework and research.

The comparisons between Research and Other respondents on the syllabus topics suggest that, in at least some cases, instructors at non-Doctoral-Extensive universities are spending more time teaching topics that seem to be less prevalent in current research[1]. Three topics taught significantly more by Other respondents are Genetic Algorithms ($p = .033$), Expert Systems ($p = .046$), and Fuzzy Logic ($p < .001$). In contrast, two topics of more recent research interest, Reinforcement Learning ($p = .005$) and the fundamental technique of Information Theory ($p = .019$) are covered significantly less by Other respondents. However, some

[1]We do not have an objective measure of "prevalence in current research," but if such a thing existed, it could be correlated with the results in Table 2. We should also note that the curriculum at some institutions may vary such that all prevalent topics are covered in several different courses, not in just the one covered by this survey.

		Total Institutions	Surveyed Institutions	Invitations Sent	Surveys Received	Response Rate
Research	**Doctoral-Extensive**	151	126	354	64	18%
Other	**Doctoral-Intensive**	112	54	84	19	23%
	Masters	611 (61)	23	29	10	34%
	Baccalaureate	550 (55)	6	6	5	83%
	Technical	25	6	6	1	17%
	N/R				3	
	Total	404	215	479	102	21%

Table 1: Number of invitations sent and surveys received, broken down by Carnegie Classification. Masters and Baccalaureate institutions were sampled at a rate of 10%. "Research" responses include the Doctoral-Extensive (Research I) respondents, while "Other" responses includes the other five categories, including Doctoral-Intensive (Research II).

currently popular areas of research, such as Bayesian Networks, are not taught at significantly different rates ($p >> .05$). These results suggest that, at least to some extent, there may be an effect of "curriculum inertia," in which topics that were formerly important to AI are still taught, despite their limited importance to current research, particularly at small, less research-oriented institutions.

One interesting note is that although Philosophy of Mind is discussed at all by only about half of all respondents, 12 respondents from Research universities (19%) said that they spend more than one hour on this topic, while none (0%) of the Other respondents spend more than an hour ($p = .005$). We have no explanation for this surprising result.

Teaching Techniques

We asked respondents a number of questions about teaching techniques in an attempt to understand how ItAI classes are being taught. We identified a number of common teaching techniques, including several that are considered aspects of cooperative pedagogy (Foyle 1995; Antony & Boatsman 1994). Cooperative (also collaborative) teaching techniques include group projects, small group activities, and other activities that involve students more actively than do traditional lectures. Instructors were asked whether they currently use the technique when teaching ItAI, have used it in the past for ItAI, have used it ever for any class, or never used that technique, choosing the first of these responses that apply. The responses are shown in Table 3. Lectures were predominant, while class discussions, group projects, and mailing lists/web discussion boards were currently being used by half or more of the respondents.

We also asked four opinion questions related to teaching techniques. Results are shown in Table 4. Note that more than 80% of respondents supported the use of group projects, but only 62% claim to have ever used them for ItAI, and only 51% currently use group projects (Table 3). In general, respondents seem to agree that interactive techniques such as class discussions and on-line discussion forums are valuable, and that group projects are also valuable. Respondents are split as to whether ItAI in particular would benefit from creative teaching techniques.

Respondents were then asked to note any teaching techniques they may have used that were either particularly successful or particularly unsuccessful for ItAI. The open-ended results were categorized and sorted by frequency, and the most common responses are shown in Table 5.

Group projects were strongly supported, but some concerns were expressed. It was suggested that faculty be careful not to over-use them, as they can leave struggling students behind, particularly if projects are cumulative over the semester.

Respondents said that students were "bored" and "not engage[d]" by lectures alone. There was support for using class discussions, as well as in-class puzzles, games, and group problem-solving of various sorts to complement lectures, as well as in-class quizzes to reinforce major themes and evaluate progress.

In other comments, current research topics seem worth discussing, but having students actually read the primary literature was not found to be valuable. The existing AI texts were criticized by a few respondents, and in particular the homework at the ends of the chapters.

Other notable comments included endorsements of visualization tools and the use of concrete examples and demonstrations for new concepts, and recommendations against an excessively-broad syllabus.

Additional Findings

The results of the survey provide a snapshot of the current state of AI instruction. In addition to the results presented in detail above, other results from the survey proved interesting as well.

We asked who is taking the class, and who is teaching it. Demographic results were were not surprising, but other information from these sections was notable. A majority of instructors reported that they were well-prepared to teach, and that their students were prepared to learn, but there were significant exceptions. About one-third of instructors said that their students lacked adequate background knowledge, and about 40% said that they themselves did not feel prepared the first time that they taught the class. It should be noted that we asked about teaching experience and awards, but there were no significant differences in syllabus or teaching techniques based on these factors.

We asked about programming and projects. Instructors strongly like programming assignments, although non-

Topic	**M**	**s²**		**Do Not Discuss**	**< 1 Hour**	**> 1 Hour**
Heuristics	2.77	.36	Research	2%	0%	98%
			Other	13%	5%	82%
1st-Order Logic	2.65	.47	Overall	12%	11%	77%
Inference	2.63	.53	Overall	15%	8%	78%
State-Space Search	2.62	.52	Research	3%	6%	91%
			Other	19%	11%	70%
Games	2.56	.57	Overall	16%	12%	72%
Machine Learning	2.56	.57	Overall	16%	12%	72%
Expert Systems	2.41	.63	Research	22%	36%	42%
			Other	18%	16%	66%
Neural Networks	2.41	.72	Overall	24%	12%	64%
Planning	2.33	.68	Overall	22%	22%	56%
0th-Order Logic	2.27	.60	Research	8%	23%	69%
			Other	24%	38%	38%
Constraint Satisfaction	2.24	.59	Overall	20%	35%	44%
Logic Programming	2.20	.70	Overall	27%	27%	47%
Probability Theory	2.20	.54	Overall	20%	42%	39%
Genetic Algorithms	2.16	.70	Research	33%	41%	27%
			Other	27%	22%	51%
Bayesian Networks	2.13	.69	Overall	29%	30%	42%
Natural Language Processing	2.10	.72	Overall	31%	27%	41%
Theorem Proving	2.02	.59	Overall	28%	41%	30%
Turing Test	1.97	.31	Overall	17%	70%	14%
Fuzzy Logic	1.96	.71	Research	48%	45%	6%
			Other	34%	24%	42%
Unsupervised Learning	1.92	.67	Overall	38%	33%	29%
Semantic Networks	1.91	.74	Overall	42%	26%	33%
Robot Navigation	1.84	.51	Overall	35%	46%	19%
Computational Learning Theory	1.75	.70	Overall	50%	25%	25%
Non-monotonic Reasoning	1.74	.46	Overall	39%	48%	13%
Distributed AI	1.73	.56	Overall	45%	37%	18%
Machine Vision	1.70	.61	Overall	50%	30%	20%
Explanation-Based Learning	1.68	.53	Overall	47%	38%	15%
Case-based Reasoning	1.65	.57	Research	66%	28%	6%
			Other	46%	32%	22%
Parsing	1.61	.46	Overall	50%	39%	11%
Philosophy of Mind	1.59	.36	Research	52%	30%	19%
			Other	46%	54%	0%
Robot Control	1.55	.42	Research	66%	20%	14%
			Other	49%	46%	5%
Reinforcement Learning	1.55	.52	Research	36%	39%	25%
			Other	68%	24%	8%
Consciousness	1.45	.38	Overall	61%	33%	6%
Robot Architecture	1.42	.31	Overall	62%	35%	3%
Information Theory	1.42	.37	Research	48%	48%	3%
			Other	71%	21%	8%
Speech Recognition	1.41	.36	Overall	65%	29%	6%
Line Labeling	1.40	.41	Research	83%	17%	0%
			Other	62%	27%	11%
Statistical NLP	1.22	.22	Overall	81%	17%	3%

Table 2: Responses for syllabus topics, sorted by mean amount taught (Do Not Discuss = 1, $<$ 1 Hour = 2, $>$ 1 Hour = 3), with weighted response rate for each option. Research and Other responses are shown separately if differences are significant ($p \leq .05$).

	Currently Used for ItAI	Ever Used for ItAI	Ever Used for Any Class	Never Used
Lectures	94%	3%	1%	2%
Class Discussions	80%	7%	5%	8%
Group Projects	51%	11%	21%	18%
Mailing List/Web Board	50%	8%	9%	33%
In-class Indiv. Exercises	33%	7%	21%	39%
In-class Group Exercises	32%	9%	20%	39%
Student Presentations	26%	17%	28%	19%
Class Surveys	18%	19%	11%	62%

Table 3: Responses for common teaching techniques. Not weighted.

Lectures are sufficient for teaching ItAI			
16%	55%	27%	2%
Strongly Disagree	Disagree	Agree	Strongly Agree
Group projects improve students' comprehension of difficult AI concepts			
0%	19%	67%	14%
Strongly Disagree	Disagree	Agree	Strongly Agree
Class discussion improves students' comprehension of difficult AI concepts			
0%	6%	67%	27%
Strongly Disagree	Disagree	Agree	Strongly Agree
Compared to other classes, ItAI can benefit more from the use of creative teaching techniques			
5%	47%	37%	11%
Strongly Disagree	Disagree	Agree	Strongly Agree

Table 4: Responses to four opinion questions on teaching techniques. Not weighted.

Beneficial	**Not Beneficial**
Large (group) projects	Lectures
Class discussions	AI texts & their homeworks
Puzzles and games	Literature survey/reading
In-class problem solving	Group/cumulative projects
Discuss open research	Too-broad syllabus

Table 5: Frequent results of open-ended questions regarding beneficial and not beneficial teaching techniques, sorted from most-frequent to less-frequent responses.

programming homeworks are popular as well. Projects, and particularly reasonably-scoped group projects, are seen as generally valuable.

We asked about textbooks, and also asked about the use of central organizing principles. A substantial majority of respondents use the AI: A Modern Approach text (Russell & Norvig 1995) (1st edition, at the time of this survey), and those respondents were somewhat happier with their text than were other respondents. To some extent, this may be due to the use of that text's agent-centric central organizing principle. More than 50% of the respondents valued agent-centrism as a method to avoid the so-called "smorgasbord" problem of AI syllabi, while more than 75% endorsed some sort of central organizing principle. Regardless, all respondents tended to think that the selection of AI texts is sub-par compared with other subjects and fields.

Conclusions

The results presented in detail above are focused on the problem of syllabus inertia and the use of cooperative teaching techniques. The results on the syllabus questions show a notable amount of variability, and in particular, showed a significant amount of "curriculum inertia," particularly at the smaller, less research-oriented institutions. The results of the teaching techniques section show that many instructors supplement lectures, particularly with class discussions and other sorts of in-class activities, so as to most successfully engage the students in learning.

This survey is one one of several empirical studies that could, and should, be performed. A complete picture of the state of AI education would not be complete without student surveys and in-class experimental comparisons. We hesitate to make too many recommendations from this data, but several suggestions do seem worth making. First, it is important to make sure that the syllabus reflects current notions of what is important in AI. Students who may have done well in an AI class with a 20-year-old syllabus could easily be unprepared for more advanced coursework. This recommendation follows from the results showing significant curriculum inertia, as well as from opinions expressed at the AAAI workshop (Hearst 1994). Second, the use of programming assignments that reinforce concepts taught in class, and the use of larger group projects, seem valuable. The data presented here show strong support for these approaches, again in accordance with the AAAI workshop recommendations. And third, as revealed by the open-ended and opinion responses regarding teaching techniques, lectures can beneficially be supplemented with in-class discussions and other cooperative pedagogical tools.

Acknowledgments

We wish to thank for their contributions and assistance the UIUC Office of Instructional Resources, Gary Dell, Wolfgang Viechtbauer, Michelle Hinn, everyone else who gave suggestions, and the participants.

References

Antony, J., and Boatsman, K. C. 1994. Defining the teaching-learning function in terms of cooperative pedagogy: An empirical study of faculty practices. In *Proceedings of the 19th Annual Meeting of the Association for the Study of Higher Education*.

Dillman, D. A.; Tortora, R. D.; and Bowker, D. 1998. Principles for constructing web surveys. SESRC Technical Report 98-50, Washington State University.

Engel, G., and Roberts, E. 2001. ACM/IEEE CS Joint Curriculum Task Force, Computing Curricula 2001, Computer Science Volume. Available at http://www.acm.org/sigcse/cc2001/.

Foyle, H. C. 1995. *Interactive learning in the higher education classroom: cooperative, collaborative, and active learning strategies*. Washington, D.C.: National Education Association.

Harris, H. D., and Kiefer, S. M. 2002. A survey of instructors of introductory artificial intelligence. UIUC Department of Computer Science Technical Report UIUCDCS-R-2002-2313.

Hearst, M. A. 1994. Preface: Improving instruction of introductory AI. In *Improving Instruction of Introductory Artificial Intelligence*. AAAI Press. Papers from the 1994 AAAI Fall Symposium, AAAI Technical Report FS-94-05.

McCormick, A. C. 2001. The Carnegie classification of institutions of higher education, 2000 edition.

Russell, S., and Norvig, P. 1994. A modern, agent-oriented approach to introductory artificial intelligence. In Hearst, M. A., ed., *Improving Instruction of Introductory Artificial Intelligence*. AAAI Press. 15–18. Papers from the 1994 AAAI Fall Symposium, AAAI Technical Report FS-94-05.

Russell, S., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice Hall.