

Department of Computer Science
CSC 444
Fall 2009

Project Guidelines

2 pp abstract, with references, due in class, Tues. Nov. 17

Project due: 5pm, Mon. December 7

The project should be a small research undertaking, with “research” broadly understood; the emphasis could be more on researching and surveying the literature than in contributing your own ideas, or it could be on pursuing some ideas of your own, presenting some results of this effort, and proposing further research. You might also put the primary emphasis on developing an interesting program. AI is partly an engineering enterprise! The topic should tie in with the subject matter of CSC 444, i.e., it should have something to do with logical approaches to an AI problem. By “logical” I generally mean quantified logic, though this stipulation might be relaxed for projects involving Bayes nets or SAT-solving, though even there some similar but more “logical” options are available, such as Markov logic networks, or the use of SAT-solving to solve problems in planning, quantified reasoning, action logic, etc.

But whichever way you slant your project, there should be evidence both of delving into the literature (with more breadth in the case of a project that is primarily a literature survey), and of *having given your topic some independent critical thought and having developed your own perspective on what the main achievements and important directions for further work are*. As some sort of guide about how much reading is expected, you should probably study about 2-4 articles carefully (explaining the gist of these in your own words, and highlighting strengths and shortcomings). Articles might be journal articles, book chapters or conference articles, though the latter tend to be too sketchy to enable a real technical grasp of the subject matter.

The project report should be around 3000-4000 words in length (though there’s no penalty for overruns). The exposition of relevant literature should be as *specific* as possible. *Particular examples are usually more informative than high-level descriptions*. You should write your report as if you were using it as lecture notes for one or two lectures to a class like CSC 244/444 on your chosen topic, making sure that your listeners/readers would acquire sufficient knowledge about your topic to be able to answer some technical questions about it on an exam.

Be sure to make clear what ideas and opinions you present are your own, and which ones are taken from your sources (which of course should be properly cited). For any programs you develop, be sure to motivate and explain your approach (in relation to your readings), to document the code clearly, to show ample examples of what the program does, to discuss its capabilities and limitations and what has been learned.

In all cases, provide a concluding section that summarizes your perspective on the work studied and carried out, and on promising directions for further work.

In accord with the above guidelines, grading will be based on

- *effort* (breadth and depth of study and research demonstrated by the essay/project),
- *cogency* (quality of organization and writing, clarity and persuasiveness of the presentation, quality and performance of the programs, if any), and
- *originality* (independence of viewpoint, critical judgement, insight, ability to formulate and develop ideas).

Sources

Good sources for your readings are

- The Brachman & Levesque text; browse the book to get an advance idea of what might interest you; look especially at further readings suggested at the end of chapters, and at material not scheduled to be covered in the course.
- The Russell and Norvig supplementary text; this covers quite a few topics not covered in B & L.
- The *AI Journal*; this is a high-quality journal with detailed technical papers in the full range of AI topics, but with emphasis on representation and reasoning. Tables of contents, abstracts, and many articles can be obtained on-line: see the web at URL
http://www.elsevier.com/wps/find/journaldescription.cws_home/505601/description#description
- *JAIR* - the *Journal of AI Research*; this is a high-quality fully on-line journal – a new and helpful trend in journals! The URL for web access is <http://www.jair.org/>
- The journal *Cognitive Science*; this contains mostly articles concerned with cognitive modelling (by neural nets or more conventional computational techniques), learning, tutoring systems, etc., but also some more “mainstream AI” articles. See <http://www.cognitivesciencesociety.org/about.html>
- the Canadian journal *Computational Intelligence*, which is general in scope and generally of very good quality. Web page at <http://www.blackwellpublishing.com/journal.asp?ref=0824-7935> .
- The *KR'89-KR'04* Proceedings; the full title is the *Proc. of the 1st - 9th Int. Conf. on Principles of Knowledge Representation and Reasoning*; this is the premier conference in the KR&R field, and the papers are generally detailed enough to be useful. See the web at <http://www.kr.org/> for information on upcoming and past conferences, including calls for papers and the list of talks at past conferences.
- The *AAAI* Proceedings; full name: the *Proc. of the Nat. Conf. on Artificial Intelligence* (sponsored by the AAAI - the Am. Assoc. for AI). This is the main US conference on AI, held annually, and covering the full AI spectrum. This gives you a very good overview of what's happening on the research frontiers; the articles are sometimes too concise to give you adequate technical depth. The AAAI home page is at <http://www.aaai.org> and the contents of past conference proceedings at <http://www.aaai.org/Press/Proceedings/proceedings.html>.

- The *IJCAI* Proceedings; full name: the *Proc. of the Int. Joint Conf. on Artificial Intelligence*. This is the main international AI conference, held in odd-numbered years since 1969. It's the oldest regular conference in AI, and is comparable in quality and scope to AAAI conferences. The IJCAI home page is at <http://www.ijcai.org/>. Last time I looked, proceedings from IJCAI-69 to IJCAI-09 were available!
- The *CSCSI* and *CAIAC* Proceedings; full name: *Canadian Society for Computational Studies of Intelligence*, renamed *Canadian Artificial Intelligence Association*. The conference is held in even-numbered years. Web page: <http://cscsi.org/>. This used to provide access to the programs of CSCSI-96 to CSCSI-0?, and ordering info for past conferences, but some links seem to have vanished. These conferences are smaller than AAAI or IJCAI, but generally of good quality.
- Look at <http://www.aaai.org/Resources/ai-organizations.html> for some further AI-oriented organizations (and hence further publications), such as the Eur. Coordinating Comm. on AI, which sponsors the European AI Conferences (ECAI), where further high quality papers may be found; the IEEE Computer Society and their AI-oriented journals; and others.
- A wiki page was created for SIGART in Dec. 2007. You can reach it directly at http://wiki.acm.org/sigart/index.php?title=Main_Page. You can find a link from the SIGART page at <http://www.sigart.acm.org/>
- Some “classic” papers are to be found in various “Readings” collections: Brachman & Levesque, *Readings in Knowledge Representation* (Morgan Kaufmann, 1985); Webber & Nilsson, *Readings in Artificial Intelligence* (Tioga, 1981); Collins & Smith, *Readings in Cognitive Science* (Morgan Kaufmann, 1988); Allen, Hendler & Tate, *Readings in Planning* (Morgan Kaufmann, 1990); Ginsberg, *Readings in Nonmonotonic Reasoning* (Morgan Kaufmann, 1981).
- Google has become a nearly universal tool for finding information on any topic, and often leads you to references you might easily have missed. (Chances are that in doing a google literature search you'll also be led to citeseer (<http://citeseer.nj.nec.com/cs>), which is a broad-coverage index of the scientific literature.) *But beware:* google will often lead you to personal and course pages of variable quality. While these can be very useful, your references should primarily be from the *refereed* literature, which is much more reliable. Also, keep in mind that you cannot lift passages from references or sites you locate and insert them in your essay as if they were your own words. This would be *plagiarism*, a serious violation of U of R's academic honesty policy. Express everything you want to say in your own words – even if you feel one of the references says it better (though of course explicit quotation from a book, journal or conference paper, with a citation, is perfectly ok).

Some possible topics

The following are some miscellaneous topics that happen to interest me; so these are topics with which I'm in a better position to help than some others. But don't regard this list as anything but a small sampling of possibilities! Also, the references mentioned are just the ones I've happened to come across – don't assume that these are at all complete or up to date! By the way, overlap between CSC 400 projects and CSC 444

projects is ok, though the two reports should emphasize different aspects of a larger project, and the total work done should be significantly more than for a 400-only or 444-only project.

1. “Natural language inference”. This is a type of reasoning that is closely based on the structure of natural language, making particular use of the properties of quantification and of polarity (upward and downward entailing environments) in natural language, as well as factivity of certain verbs (e.g., *discover that ...*). I’d be very interested in seeing someone report on this, particularly since our own EPILOG system uses a natural logic – one close to ordinary language. Bill MacCartney, working with Chris Manning at Stanford, has recently completed his thesis on this topic, and had some excellent previous papers with Manning (including the “best paper” at COLING’08). The thesis is

Bill MacCartney, 2009, *Natural language inference*, Ph.D. dissertation, Stanford University, June, <http://nlp.stanford.edu/~wcmac/>.

2. CLib – the Component Library. This refers to an ontology and basic conceptual knowledge base, containing such knowledge as that when x enters y , it does so through some “portal” z , and is first outside of y , and after entering is inside of y . I think the first general description of this was given in

P. Clark and B. Porter, “Building concept representations from reusable components”, in *Proc. of the 14th Nat. Conf. on Artificial Intelligence (AAAI-97)*, pp. 369-376, AAAI Press, CA, 1997.

and the project has continued under the direction of Bruce Porter (at U.Texas, Austin). A general description and assessment would be interesting to see – even more so if it turned out to be possible to acquire and experiment with this system (I think it is, and I could help with contacting Porter.)

3. Commonsense psychology. Make a study of the work that has been done by Jerry Hobbs and Andrew Gordon in formalizing our intuitive understanding of notions like knowing, remembering, finding out, forgetting, etc. One publication is

J. R. Hobbs, A. S. Gordon, “Toward a large-scale formal theory of commonsense psychology for metacognition”. *2005 AAAI Spring Symp. on Metacognition in Computation*, AAAI Tech. Rep. SS-05-04, Stanford Univ., pp. 49-54, March 21-23, 2005.

but there are more recent ones as well. Andrew Gordon has said he’d be happy to supply the axiom base to us, including with English verbalizations (for better understanding) that he would produce upon request.

4. Causal net (Bayes net) inference based on “quasi-probabilities” – a new characterization of probabilities in such networks. Ref:

<http://www.cs.rochester.edu/~schubert/papers/uai04.pdf> These quasi-probabilities are manipulated much like ordinary algebraic expressions based on $+$, $-$, and product, except that ordinary product is replaced by a “weak product” operator $*$ that behaves a bit like set union and is idempotent. For example, $p * p = p$, $(p * q * r) * (q * r * t) = p * q * r * t$, $p * (1 - p) = p - p = 0$, etc. It turns out that using quasi-probabilities, we can compute joint probabilities for any set of nodes as the (weak) product of their individual quasi-probabilities. This gives a

theoretically very straightforward new way of computing probabilities, including arbitrary conditional probabilities in “noisy AND/OR/NOT” causal nets – and in fact all Bayes nets with true/false nodes can be recast as such nets.

An undergraduate, David Sloan, implemented some nice programs for quasi-probability manipulation, but this needs to be taken further, the most interesting possibility being the use of approximations that eliminate low-probability terms, as a way of gaining speed. Another undergraduate, Tom Weingarten, implemented a square-wave analogue representation of symbolic quasi-probabilities (somewhat reminiscent of neural nets), and demonstrated the feasibility of using this for numerical calculation of Bayes-net probabilities for small networks. There are also intriguing possibilities for generalizing to quantified logics (see also next item).

This would be suitable for someone with prior interest in and knowledge of Bayes nets. There are some vaguely relevant papers in the literature (D’Ambrosio, Symbolic probabilistic inference in belief nets, 1989, Shachter, D’Ambrosio & Del Favero, AAAI-90, Darwiche & Pearl, AAAI-94, Gerard Vreeswijk, 1999, Lang, Liberatore, & Marquis, 2002)

5. Probabilistic planning. Survey some recent work in this area, e.g., in hierarchical probabilistic planning, or in *knowledge-rich* probabilistic planning (most past work has considered planning only in very narrowly defined worlds, rather than worlds where large amounts of knowledge would be relevant to, and play a role in, planning).

Another possible angle is to look at the possibility of applying the algebraic probability logic mentioned above, as described in the UAI 2004 paper, to forward projection of probabilities in plans, as a way of determining probabilities of various properties holding at various times. This would be for someone who is quite mathematically inclined.

6. Other recent work in planning, including SAT-planning, fast-forward planning, plan abstraction, etc. SAT-planning is scheduled to be covered briefly in the course, but you could branch out from there to related work, or altogether different topics. Some themes and associated names that come to mind are (as mentioned) SAT-planning (e.g., Kautz & Selman, 1996, 1999, Ernst, Millstein & Weld, 1997), planning using integer programming (e.g., Bockmayr & Dimopoulos 1999, Vossen, Ball, Lotem, & Nau, 1999), the Graphplan approach (Blum & Furst, 1995, Koehler, Nebel, Hoffman & Dimopoulos, 1997, Long & Fox 1999, Gerevini & Serina, 2002), forward search (Bacchus & Kabanza, 2000, Jörg & Nebel, 2001, Refanides & Vlahavas, 2001), or methods employing state-space search (e.g., Bonet & Geffner, 1999, 2001). An interesting fairly recent effort in automatic abstraction in planning is that of Botea, Enzenberger, Müller, & Schaeffer (2005), Botea (2006), and Botea, Müller, & Schaeffer (2007). The planning competitions at the AIPS and ICAPS conferences give a state-of-the-art picture of the latest planning technology.
7. Survey and evaluate recent work on “continual planning”. See the special issue of AI Magazine, vol. 20, no. 12 (winter 1999); also, for example, “Integrated planning and execution for autonomous spacecraft,” by Chien *et al.* in *Proc. of the IEEE Aerospace Conference (IAC)*. But find more recent work as well.
8. Discovering “macro-operators” for accomplishing goals in planning domains by analyzing properties of, or experimenting with, various sequences of two or more

operators applied in succession. There was some early work on ‘MACROPS’ by Nils Nilsson, and some more recent, possibly relevant work on discovering plan operator properties (see K.S. Tae, D.J. Cook, and L.B. Holder, “Experimentation-driven knowledge acquisition for planning”, *Computational Intelligence* 15(3), pp. 252-279, 1999). Also there has been work on finding useful sequences of operations in domains like Rubik’s cube, using an algebraic perspective (Richard Korf may be a relevant name here). No doubt there is other work, waiting to be found by anyone who makes a more serious search. Daphne Liu could help find references in this area, I believe.

9. Concept formation and discovery in problem solving. When people solve non-trivial problems, they spend a good deal of time *exploring* the problem domain, forming helpful concepts and establishing general properties of the domain. This is a very important research area if we want to boost the problem-solving ability of computers. Some possibly relevant references:

M. Sims, “IL: An AI approach to theory formation in mathematics”, Tech. Rep. ML-TR-33, Dept. of Computer Science, Rutgers Univ., 1990.

S. Epstein, “On the discovery of mathematical theorems”, IJCAI ’87.

S. Epstein and N. Sridharan “Knowledge representation for mathematical discovery: three experiments in graph theory”, *Applied Intelligence* 1(1), 7-33, 1991.

R. Davis and D. Lenat, Knowledge-Based Systems in Artificial Intelligence, McGraw-Hill, 1982.

D. Lenat, “EURISKO: ...”, in *Artificial Intelligence* 21, 1983.

K. Haase, Exploration and Invention in Discovery, Ph.D. Thesis, MIT, 1990.

K. Haase, “Cyrano-3: An Experiment in Representational Invention”, Workshop on Machine Learning, *Int. Conf. on Machine Learning*, 1992.

A. Bundy, S. Colton, and T. Walsh (1998), “HR - Automatic concept formation in finite group theory”, to appear in the Int. Congress on Discovery and Creativity, Gent, Belgium. [must be available by now]

(See also Colton’s PhD dissertation, “HR - Automatic Concept Formation in Finite Algebras”, PhD dissertation, Dept. of AI, U. Edinburgh, 80 South Bridge, Edinburgh, EH1 1 HN; Scotland, simonco@dai.ed.ac.uk.)

T. Recio and M.P. Vélez, “Automatic discovery of theorems in elementary geometry”, *J. of Automated Reasoning* 23(1), July 1999, pp. 63-82.

S. Colton, A. Bundy, and T. Walsh, “Automatic invention of integer sequences”, *Proc. AAAI-2000/IAAI-2000*, pp. 558-563.

S. Colton, A. Bundy, and T. Walsh, “HR: Automatic concept formation in pure mathematics”, *Proc. 16th IJCAI*, 1999.

S. Colton, A. Bundy, and T. Walsh, “Automatic identification of mathematical concepts”, *Proc. 17th Int. Conf. on Machine Learning*, 2000.

Peter Gaerdenfors, *Conceptual Spaces*, MIT Press 2000. (Relevant to concept formation, using geometrical models.)

C. Larson, “Intelligent machinery and discovery in mathematics”. Unpublished preprint, available from <http://math.uh.edu/clarson/>, 1999.

F.J. Costello and M.T. Keane, "Efficient creativity: constraint-guided conceptual combination", *Cog. Sci.* 24(2), Apr-Jun 2000, 299-??? (end of issue).

- Reasoning (seemingly) dependent on induction. David McAllester (formerly of MIT, now chief of Toyota Tech. Inst.) has worked on a system ("ONTIC") for such problems, with applications to program verification in mind. He gives some examples that are simple for people but hard for machines:

"A king on an empty chess board can reach every square. A pawn can not reach the first rank without first becoming a queen. If I repeatedly remove objects from a box, the box will eventually become empty. Consider a "rotation" operation on a row of marbles where I take the first marble and put it at the end. If I rotate the row as many times as there are marbles they end up in the same order they started in. All of these examples are challenging for automated systems. For example, the rotation example can not be solved by the Boyer-Moore prover without a human providing a non-obvious induction invariant on the rotation operation."

D. McAllester, "Observations on cognitive judgements", *Proc. 9th Nat. Conf. on Artificial Intelligence (AAAI-91)*, July 14-19, 1991, MIT Press, Cambridge, MA, pp. 910-914.

His 1995 talk, "Obviousness: A Window on the Mind?" (Slides from a talk at <http://www.research.att.com/~dmac/>) gives additional interesting examples, some of them aimed against imagistic approaches.

J. of Automated Reasoning, special issue on Automated Mathematical Induction, ed. Hantao Zhang, 16(1-2), 1996.

- Theorem proving by model elimination (and related methods). This class of approaches has led to some visible successes in the last decade. One recent prominent system is E-SETHEO. A list of systems entered in the CADE-13 theorem proving competition can be found at <http://www.cade-19.info/> (click on "Competition", and then on "Systems and entrants). E-SETHEO was a past winner. Information on E-SETHEO can be found at

<http://www4.informatik.tu-muenchen.de/schulz/WORK/e-setheo.html>

A paper about efficiency gains through *lemma* generation in a preprocessing phase:

Marc Fuchs, "System description: similarity-based lemma generation for model elimination", in 15th Int. Conf. on Automated Deduction, vol 1421 of *Lecture Notes in Computer Science*, pp. 33-37, 1998.

- Augmenting general reasoning with "imagistic" reasoning about motion of complex deformable objects. How do we know that a bunched-up sock can fit into a coffee cup, but not a bunched-up sweater? That a closed umbrella can easily be put in the trunk of a car, but an open one with difficulty or not at all? That a wet dishcloth makes a better flyswatter than a silk kerchief? Neither logical inference, nor computational geometry, seem like quite the right tools for realizing such reasoning in machines. I have some sketchy ideas about how one might proceed, and some possibly relevant references (check out the first one, especially, which is very recent):

B. Johnston and M.-A. Williams, "Autonomous learning of commonsense simulations", *Commonsense 2009*, Toronto, June 1-3.

Janice Glasgow, *Diagrammatic Reasoning: Cognitive & Computational Perspectives*, AAAI Press Series (Paperback), M I T Press, 1995.

T. Dar, L. Joskowicz, "Understanding mechanical motion: From images to behaviors", *Artificial Intelligence* 112(1-2), 1999, pp 147-179. T Dar, L Joskowicz

Nigel J.T. Thomas, "Are theories of imagery theories of imagination? An active perception approach to conscious mental content", *Cognitive Science* 23(2), Apr-Jun 1999, 207- ??.

1998 AAAI Fall Symposium on Formalizing Reasoning with Visual and Diagrammatic Representations, Oct. 23-25, Orlando, FL (see www.aaai.org/Symposia/Fall/). The meeting was chaired by Gerard Allwein (Indiana U.) and Kim Marriott (Monash U.).