

Department of Computer Science
CSC 444
Fall 2011

Project Guidelines

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

Project due before or in class, Tues. December 6

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 3-5 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 3500-4500 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/journal_csj.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.php> .

- 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-11 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 many links seem to have vanished. The 2009 conference proceedings are said to be available in the Lecture Notes in Artificial Intelligence series of Springer Verlag; maybe others as well. These conferences are smaller than AAAI or IJCAI, but generally of good quality.
- Look at <http://www.aiinternational.org/societies.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://citeseerx.ist.psu.edu/index.jsp>), 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, conference paper, or even website, 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. *Solving FOL challenge problems in EPILOG*. This would be an interesting programming problem with AI research & development significance. EPILOG is a reasoning engine using a natural-language-like knowledge representation, Episodic Logic (EL). EPILOG 2, the most recent version, is locally available at </u/epilog/EPILOG2/epi2> .

A quick usage overview is given at </u/epilog/EPILOG2/README>. The documents in </u/epilog/EPILOG2/docs/> are for the earlier version of EPILOG, but most of the descriptions apply to EPILOG 2. (However, EPILOG 2 lacks forward inference and the specialist reasoners of the earlier version. These are not needed for FOL theorem proving.)

EL subsumes FOL, but also allows for operations like generalized quantification (“*most people who like dogs*”), modification (“*very nearly crashed into a tree*”), and reification (“*the fact that Armstrong walked on the Moon*”). It has often been claimed that such expressive power makes efficient reasoning impossible, but Fabrizio Morbini, the designer of EPILOG 2 (the most recent version) showed in 2009 that EPILOG can hold its own in FOL theorem proving against such highly optimized FOL theorem provers as VAMPIRE:

F. Morbini and L.K. Schubert, “Evaluation of Epilog: A reasoner for Episodic Logic”, *Commonsense'09*, June 1-3, 2009, Toronto, Canada. Also available at <http://commonsensereasoning.org/2009/papers.html> .

However, this work is somewhat out of date, because new theorem proving competitions, using updated challenge sets, take place every year. The challenge sets are maintained in the TPTP (“thousands of problems for theorem provers”) Problem Library, by Geoff Sutcliffe at the Univ. of Miami, at

<http://www.tptp.org>, with full documentation at <http://www.tptp.org/TPTP/TR/TPTPTR.shtml> .

The problem in testing EPILOG against these challenge sets is that the syntax and format in which they are stored must be converted to the form required by EPILOG, and EPILOG must then be run on the chosen problem sets in an automated fashion, collecting timing and success/failure statistics. Fabrizio Morbini, who has programmed such tests previously (and now works at the Institute for Creative Technologies) would be available for help and advice in such a project. (I can supply his email address.) Note that Morbini's programs for running TPTP problems and VAMPIRE (which still seems to be the strongest competitor) are in </u/epilog/EPILOG2/epi2/> .

2. “*Natural Logic (NLog)*” (*esp. in the context of “textual inference*”). This is a type of reasoning that is closely based on the structure of natural language. For example, from the premise “*I know that you won't forget to give me a call*”, NLog allows immediate inference of “*You will contact me*”, based on the lexical facts that knowing something implies that it is true, that forgetting to do something

implies not doing it (and this is reversed under negation), and giving someone a call implies contacting them. Besides exploiting such lexical facts, NLog makes use of the notion of positive and negative environments (“polarity”) and “upward and downward entailment” properties of quantifiers, connectives, and other types of words. I’d be very interested in seeing someone dig into the recent literature on this, particularly since our own EPILOG system uses a representation very close to ordinary language, and polarity-based reasoning close to NLog methods. Bill MacCartney, working with Chris Manning at Stanford, completed his thesis on this topic in 2009, 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/>.

We have begun exploring connections more closely in our work here, e.g.,

L.K. Schubert, B. Van Durme, and M. Bazrafshan, “Entailment inference in a natural logic-like general reasoner”, *Common Sense Knowledge Symposium (CSK-10)*, AAAI 2010 Fall Symposium Series, Arlington, VA, November 11-13, 2010; and

K. Stratos, L.K. Schubert, & J. Gordon, “Episodic Logic: Natural Logic + reasoning”, *Int. Conf. on Knowledge Engineering and Ontology Development, (KEOD’11)*, Oct 26-29, Paris, 2011,

both available at my website <http://www.cs.rochester.edu/~schubert/>

3. *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.)

4. *From WordNet to Logic*. WordNet is the online lexicon most widely used by AI researchers (see <http://wordnetweb.princeton.edu/perl/webwn>). However, it is an informal lexicon, and if we want to use it for reasoning, we need to formalize its contents as far as possible. WordNet expresses word senses as “synsets” – sets of words that can be used synonymously in a certain sense. For example, while “head” is listed as having 33 senses as a noun, the synset $\{head, caput\}$ uniquely identifies the notion of a person’s or animals’s physical head, while e.g., $\{head, chief, top dog\}$ uniquely identifies the notion of “the person who is in charge”. (Sometimes the sets of words in a synset do not by themselves uniquely identify a particular sense, but the senses are also distinguished via their unique WordNet names; also paraphrases/glosses and examples are provided.)

WordNet provides various semantic relations between word senses (synsets), most importantly *hyponym/hypernym* relations. For example a hyponym (special case) of “head”, in the sense of $\{head, caput\}$, is “human head”, and a hypernym (gener-

alization) is “*external body part*”. (Note that short phrases like “*human head*” and “*external body part*”, not just single words, occur as WordNet lexical items.) We can put such relations into a logical format by writing, for example,

(all x ((human-head x) => (head1 x))),
(all x ((head1 x) => (external-body-part x))).

Interestingly, we can also often assume that distinct hyponyms of the same word sense are mutually exclusive. For example, “*external body part*” has various hyponyms, including of course {*head, caput*}, but also {*neck, cervix*}, {*breast, chest*}, etc. We can thus write

(all x ((head1 x) => (not (neck1 x)))), or equivalently
(all x ((not (head1 x)) or (not (neck1 x)))),

and similarly for other hyponyms of “*external body part*”, taken pairwise. As it turns out, such logical transcriptions of WordNet relations are not always “safe”. For example, problems occur because some nouns have implicit arguments, while others don’t. For example, there is something odd about saying that “*Every size is a property*” and “*every property is a concept*” (using the hypernym chain size/property/concept in WordNet), because while a size is usually thought of as the size of something (e.g., the 4000 mile radius of the Earth), a concept may be just an idea, such as the idea of a Turing machine, which just isn’t the value of anything. Also, though a “*professional*” has the hyponyms {*educator, pedagogue, pedagogue*} and {*lawyer, attorney*}, these professions are not mutually exclusive – a lawyer may also be an educator. This is a general observation about *role nouns*. These are just 2 examples of various subtleties that can cause problems in “mining” WordNet for logical relations among word senses. But it’s still the best source we have, and work in carrying out such mining to obtain a large set of lexical axioms is urgently needed. This would be very useful to our work in NLP and AI more broadly. Some previous work done here is reported in

A.N. Kaplan and L.K. Schubert, “Measuring and improving the quality of world knowledge extracted from WordNet”, Tech. Rep. 751, Dept. of Computer Science, Univ. of Rochester, Rochester, NY 14627-0226, May 2001. See

<https://urresearch.rochester.edu/institutionalPublicationPublicView.action?institutionalItemId=5>

5. *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. He is at the Institute for Creative Technologies, <http://people.ict.usc.edu/~gordon/> (where you can find more recent publications).

6. 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 described in the UAI 2004 paper,

L.K. Schubert, "A new characterization of probabilities in Bayesian networks", *Uncertainty in AI: Proc. of the 20th Conference (UAI 2004)*, Banff, Canada, July 7-11. Extended version forthcoming as tech. rep., Dept. of Comp. Sci., Univ. of Rochester, Rochester, NY 14627-0226,

(again on my web page at <http://www.cs.rochester.edu/~schubert/>). This could be used for symbolic 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.

7. 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.
8. 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.
9. 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.
10. 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.

12. 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.

Deb Roy, Kai-Yuh Hsiao, and Nikolaos Mavridis, “Mental imagery for a conversational robot”, *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(3), 2004, 1374-1383. pdf

<http://web.media.mit.edu/~dkroy/publications/index.html>

Janice Glasgow, *Diagrammatic Reasoning: Cognitive & Computational Perspectives*, AAAI Press Series (Paperback), MIT 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.).