Technical Area:
Common Sense Reasoning and Intelligent User Interfaces

Examiner: Henry Lieberman

Research Scientist, Ph.D., Media Laboratory
Massachusetts Institute of Technology

Description

One reason why software and robotic agents are helpful for humans is that they can function autonomously (at least partially). Functioning with autonomy in turn requires being context sensitive. One important aspect of context sensitivity is the ability to reason with *common sense*. Although research has been conducted in this area for years now, artificial common sense reasoning does not seem to be developed far enough to be applied to "real-world" products and applications. Nevertheless, could implementing such capabilities help create better user interfaces and autonomous agents? Could intelligent user interfaces to such entities profit from common sense reasoning methods?

Written Requirement

The written requirement for this area will consist of a 24-hour take-home exam.

Signature: _____ Date: _____

Reading list

Besides classics like *John McCarthy*, *Marvin Minsky*, *Ernest Davis*, *Mark Maybury* (intelligent user interfaces), and *Jeff Bradshaw* (agents), a large portion of the reading list consists of papers and books by people who are working on actual common sense reasoning systems, like *Doug Lenat* (CYC), *Erik Muller* (ThoughtTreasure), *Push Sing* (Openmind Common Sense), and *Tim Berners-Lee* (Semantic Web). Another chunk of readings are conference papers published by *Intelligent User Interfaces* (1999 to 2001) and *Common Sense 2001*. Finally, there are papers that describe efforts to apply some (partial) kind of common sense reasoning to robotic and software agents.

John McCarthy (1959). *Programs with Common Sense*. In Mechanisation of Thought Processes, Proceedings of the Symposium of the National Physics Laboratory, London, U.K.: Her Majesty's Stationery Office, pp. 77-84.

http://www-formal.stanford.edu/jmc/mcc59/mcc59.html
http://www-formal.stanford.edu/jmc/mcc59.pdf

Selected chapters from Ernest Davis (1990). *Representations of Commonsense Knowledge*. San Mateo, CA: Morgan Kaufmann Publishers, Inc. (book, 544 pages)

"A central goal of artificial intelligence is to give a computer program commonsense understanding of basic domains such as time, space, simple laws of nature, and simple facts about human minds. Many different systems of representation and inference have been developed for expressing such knowledge and reasoning with it. *Representations of Commonsense Knowledge* is the first thorough study of these techniques."
http://www.mkp.com/books_catalog/catalog.asp?ISBN=1-55860-033-7

Marvin Minsky (2000). *Deep issues: commonsense-based interfaces*. Communications of the ACM 43(8):66-73 (journal paper, 8 pages)

http://commonsense.media.mit.edu/minsky.pdf
http://library.navo.navy.mil/Journals/jcs03Aug00.pdf

J Jeff Bradshaw (1997). *Software Agents*. Cambridge, MA: The MIT Press (book, 490 pages)

"Recent trends have made it clear that software complexity will continue to increase dramatically in the coming decades. The dynamic and distributed nature of both data and applications require that software not merely respond to requests for information but intelligently anticipate, adapt, and actively seek ways to support users. Not only must these systems assist in coordinating tasks among humans, they must also help manage cooperation among distributed programs"
http://aaai.org/Press/Books/Bradshaw/bradshaw.html
http://aaai.org/Press/Books/Bradshaw/bradshaw-contents.pdf

Donald A. Norman (1988): *The Psychology Of Everyday Things*. New York, NY: Basic Books Inc. (book, 257 pages)

http://www.santafe.edu/~shalizi/reviews/everyday-things/
http://www.cul.co.uk/books/bcbk3.htm

Douglas B. Lenat and Ramanathan V. Guha (1990). *Building Large Knowledge-Based Systems: Representations and Inference in the CYC Project*. Reading, MA: Addison Wesley (book, 400 pages)

http://www.amazon.com/exec/obidos/ISBN%3D0201517523/002-3739191-0608065
http://www.cc.gatech.edu/~jimmyd/summaries/lenat1990-1.html

more on Cyc at:
http://www.cyc.com/publications.html
http://www.faqs.org/faqs/ai-faq/general/part4/section-26.html
http://www.ladseb.pd.cnr.it/infor/ontology/Papers/Ontobiblio/ImplOntologies.html
http://cortex.media.mit.edu/cycdoc/
http://cortex.media.mit.edu/cycdoc/readme.html

Douglas B. Lenat (1995). *Cyc: A Large-Scale Investment in Knowledge Infrastructure.* Communications of the ACM, 38(11):32-38 (journal paper, 8 pages)
http://www.acm.org/pubs/articles/journals/cacm/1995-38-11/p33-lenat/p33-lenat.pdf

Douglas B. Lenat (1998). *The dimensions of Context space.* Austin, TX: Cycorp (report, 78 pages)
"Contexts have historically been either ignored completely or else treated as indivisible atoms. A decade ago, as part of our work on building the large Cyc knowledge base of human common sense and common knowledge, our group began to study and harness the internal structure of that "atom". Each context was said to have assumptions and content; there was a theory of importing assertions across contexts; contexts were fully reified first-class terms in the CycL representation language; they were partially ordered by specialization to control visibility and access to content; and so on. Over the last few years, we've identified a finer internal structure to a context: a dozen dimensions along which contexts vary; conversely, each region of that 12-dimensional space implicitly defines a context. In effect that space is the space of assumptions, and each assertion can be thought to hold true in some region of that space. A more advanced calculus of contexts is required to handle those 12-dimensional constructs, but it enables a much more efficient, much more focused sort of "virtual lifting" of assertions from one context to another, and - by providing a superstructure that can serve as a principled guide to orient the working KB builder or peruser - makes it easier to specify the proper context in which an assertion (or question) should be stated."
http://www.cyc.com/context-space.doc

Mark J. Stefik and Stephen W. Smoliar (1993). *The Commonsense Reviews* – Eight reviews of: Douglas Lenat and Ramanathan V. Guha (1990) Building Large Knowledge-Based Systems: Representations and Inference in the CYC Project, Addison-Wesley, and Ernest Davis, Representations of Commonsense Knowledge, Morgan Kaufmann 1990. Artificial Intelligence, 61: 37-179 (journal paper, 142 pages)
Eight substantial reviews of the key book on Soar, with a response from Rosenbloom & Laird. Although the talk is cognitive, the approach is predominantly from an AI (rather than a psychological) perspective.
http://www.cs.brandeis.edu/~brendy/CYC_report.txt
http://www1.elsevier.nl/inca/publications/store/5/0/5/6/0/1/

Ramanathan V. Guha and Douglas B. Lenat (1993). *Response: Re: CycLing paper reviews.* Artificial Intelligence 61:149-174 (journal paper, 25 pages)

Ramanathan V. Guha and Douglas B. Lenat (1994). *Enabling agents to work together.* Communications of the ACM, 37(7):127-142 (journal paper, 16 pages)
http://www.acm.org/pubs/citations/journals/cacm/1994-37-7/p126-guha/
http://www.acm.org/pubs/articles/journals/cacm/1994-37-7/p126-guha/p126-guha.pdf

Push Sing (2001). *The Public Acquisition of Commonsense Knowledge*. Paper submitted to KCAP 2001 (paper, 8 pages).
http://www.media.mit.edu/~push/KCAP2001.pdf

Openmind Commonsense - Teaching computers the stuff we all know (web documents)
http://commonsense.media.mit.edu/

Erik Muller (1999): *Prospects for in-depth story understanding by computer* (unpublished paper, 23 pages)
http://www.media.mit.edu/~mueller/papers/storyund.html

Erik Muller (1999): *A database and lexicon of scripts for ThoughtTreasure* (unpublished paper, 12 pages)
http://www.signiform.com/tt/htm/script.htm

Erik Muller (2000). *A calendar with common sense*. Proceedings of the 2000 International Conference on Intelligent User Interfaces. New York, NY: Association for Computing Machinery (paper, 4 pages)
http://www.panix.com/~erik/pubs/sensical.htm
http://www.acm.org/pubs/citations/proceedings/uist/325737/p198-mueller/
http://www.acm.org/pubs/articles/proceedings/uist/325737/p198-mueller/p198-mueller.pdf

Roger Schank (1995). *Information is Surprises*. Chapter 9 from John Brockman (ed.) *The Third Culture: Beyond the Scientific Revolution*. New York, NY: Simon and Schuster (book chapter)
http://www.edge.org/documents/ThirdCulture/q-Ch.9.html
http://www.edge.org/documents/ThirdCulture/a-TC.Cover.html

Tim Berners-Lee, James Hendler, and Ora Lassila (2001). *The Semantic Web.* Scientific American Volume 284, Number 5, May 2001, pp. 34-43 (journal paper, 9 pages)
http://www.scientificamerican.com/2001/0501issue/0501berners-lee.html

Tim Berners-Lee (1998). *Semantic Web Road map* (web document, 9 pages)
http://www.w3.org/DesignIssues/Semantic.html

Tim Berners-Lee (1998). *What the Semantic Web can represent* (web document, 6 pages)
http://www.w3.org/DesignIssues/RDFnot.html

Selected papers from the *Common Sense 2001: 5th Symposium on Logical Formalizations of Commonsense Reasoning*. May 20-22, 2001
http://www.cs.nyu.edu/faculty/davise/commonsense01/

Stuart C. Shapiro, Eyal Amir, Henrik Grosskreutz, David Randell, and Mikhail Soutchanski (2001). *Commonsense and Embodied Agents: A Panel Discussion*. Common Sense 2001: 5th Symposium on Logical Formalizations of Commonsense Reasoning, May 20-22, 2001 (panel, 14 pages)

"The purpose of this panel is to discuss the implications for commonsense knowledge representation and reasoning of implementing a robot/agent with all or most of the following characteristics:

- The agent can act.
- The agent can perceive.
- The agent can communicate with humans.
- The agent does first-person reasoning: the knowledge base and reasoning are the beliefs and reasoning of the agent, rather than about it.

The agent does on-line reasoning: perceiving, reasoning and communicating with humans is performed while the agent is acting, rather than before it begins acting.

We are not distinguishing the terms \embodied agent" and \cognitive robot."
http://www.cs.nyu.edu/faculty/davise/commonsense01/final/panel.pdf

Sheila McIlraith and Tran Cao Son (2001). *Adapting Golog for Programming the Semantic Web*. Common Sense 2001: 5th Symposium on Logical Formalizations of Commonsense Reasoning, May 20-22, 2001 (paper, 8 pages)

"Motivated by the problem of automatically composing network accessible services, such as those on the World Wide Web, this paper proposes an approach to building agent technology based on the notion of generic procedures and customizing user constraint. We argue that an augmented version of the logic programming language Golog provides a natural formalism for programming Web services. To this end, we adapt and extend the Golog language to enable programs that are generic, customizable, and usable in the context of the Web. We realize these extensions in an augmented ConGolog interpreter that combines online execution of information-providing Web services with offline simulation of world-altering Web services, to determine a sequence of Web Services for subsequent execution. Our implemented system is currently interacting with services on the Web. "
http://www.cs.nyu.edu/faculty/davise/commonsense01/final/mcilraith-son2.ps

Eyal Amir and Pedrito Maynard-Reid II (2001). *LiSA: A Robot Driven by Logical Subsumption.* Common Sense 2001: 5th Symposium on Logical Formalizations of Commonsense Reasoning, May 20-22, 2001 (paper, 8 pages)

"This paper describes an implemented robot-control system that is based on Brooks-style subsumption [Brooks86] of logical theories. It implements Brooks-style subsumption between layers using nonmonotonic reasoning. We describe the control and reasoning algorithms and some of the experiments that we did with the system, running on a Nomad200 robot and a set of computers. Our experimental study shows that commonsense theories and general-purpose first-order logic theorem provers can be used to control real-time agents and robots in particular. Our system improves over traditional subsumption systems in several ways. It allows the user to send new axioms to each of the layers as the robot is running, allowing the user to give advice to the robot and to correct behaviors in runtime. Our system has no voting scheme for deciding on the behavior that should be followed. Instead, the layers work in synergy to provide the compound behavior. Our system improves over other robot-control systems that are based on logic in that it allows full first-order expressivity and that it is fully declarative."
http://www.cs.nyu.edu/faculty/davise/commonsense01/papers.html#Amir
http://www.cs.nyu.edu/faculty/davise/commonsense01/final/amir.pdf

Selected papers from the Intelligent User Interfaces (IUI) conferences 1999–2001
http://dblp.uni-trier.de/db/conf/iui/iui1999.html
http://dblp.uni-trier.de/db/conf/iui/iui2000.html
http://dblp.uni-trier.de/db/conf/iui/iui2001.html

Yuri Gawdiak, Jeffrey M. Bradshaw, Williams, B., and Hans Thomas (2000). *R2D2 in a softball: The Personal Satellite Assistant.* In H. Lieberman (ed.), Proceedings of the ACM Conference on Intelligent User Interfaces IUI 2000, pp. 125-128 (paper, 4 pages)
http://dev.acm.org/pubs/citations/proceedings/uist/325737/p125-gawdiak/

Mark Maybury (1999). *Intelligent User Interface: an Introduction*. Proceedings of the 1999 International Conference on Intelligent User Interfaces IUI '99
http://www.mitre.org/resources/centers/it/maybury/mark.html
http://www.acm.org/pubs/citations/proceedings/uist/291080/p3-maybury/
http://www.acm.org/pubs/articles/proceedings/uist/291080/p3-maybury/p3-maybury.pdf

Selected chapters from Mark T. Maybury and Wolfgang Wahlster (eds.) (1997) *Readings in Intelligent User Interfaces*. Menlo Park, CA: Morgan Kaufmann (book, 672 pages)

"This book represents a collection of the classic and contemporary readings in the field of Intelligent User Interfaces. An invaluable resource for students, professors, research scientists and engineers, it includes both fundamental research and applied innovations in the key areas of IUI including input analysis, output generation, user and discourse adapted interaction, agent-based interaction, model-based interface design, and evaluation."
http://www.mkp.com/books_catalog/catalog.asp?ISBN=1-55860-444-8

Henry Lieberman and Ted Selker (2000). *Out of context: Computer systems that adapt to, and learn from, context.* IBM Systems Journal Vol. 39, Nos. 3 & 4, pp. 617-632 (paper, 16 pages)
http://www.research.ibm.com/journal/sj/393/part1/lieberman.pdf

Jeffrey M. Bradshaw, Maarten Sierhuis, Yuri Gawdiak, Hans Thomas (2000). *Human-Centered Design for the Personal Satellite Assistant.* Unpublished paper (paper, 7 pages)
"The Personal Satellite Assistant (PSA) is a softball-sized flying robot designed to operate autonomously onboard manned spacecraft in pressurized micro-gravity environments. We describe how the Brahms multi-agent modeling and simulation environment in conjunction with a KAoS agent teamwork approach can be used to support human-centered design for the PSA."
http://citeseer.nj.nec.com/419638.html
http://ic.arc.nasa.gov/ic/publications/pdf/2000-0157.pdf

Thomas Längle, Thomas Hoeniger and Lanjuan Zhu (1997). *Cooperation in Human-Robot-Teams.* Unpublished paper (paper)
"In this paper, the concept of human-robot-team is presented. Because of its high flexibility and adaptability, the human-robot cooperation is expected to have a wide range of applications in uncertain environment not only in future construction and manufacturing industries but also for services. A Multi-Agent control architecture gives an appropriate frame for the flexibility of the human-robot-team. Robots are considered as intelligent autonomous assistants of human, which can mutually interact on a symbolic level and a physical level. This interaction is achieved with the communication between human and robots, the interpretation of the transmitted information, the coordination of the activities and the cooperation between independent agents. Equipped with some sensing modalities for the perception of the environment, the two-arm Karlsruhe Autonomous Mobile Robot (KAMRO) is introduced to demonstrate the principles of the cooperation among human and robot agents. Finally, headlines are given for future studies on human-robot-cooperation."
http://wwwipr.ira.uka.de/internal/download.php?id=976189908&filetype=pdf
http://citeseer.nj.nec.com/laengle97cooperation.html

Thomas Lüth and Jürgen Bier (1999): *Robot Assisted Intervention in Surgery*. In J.M. Gilsbach and H.S. Stiehl (eds.) *Neuronavigation – Neurosurgical and Computer Scientific Aspects*, New York, NY: Springer (book chapter)
Good survey paper about surgical robotics (but no natural language interface, and no common sense reasoning, of course)
http://www.charite.de/rv/mkg/srl/publications/GILSBACH99.pdf

Eva Stopp, Klaus-Peter Gapp, Gerd Herzog, Thomas Längle, and Tim C. Lüth (1994). *Utilizing Spatial Relations for Natural Language Access to an Autonomous Mobile Robot*. Unpublished paper (paper, 16 pages)
"Natural language, a primary communication medium for humans, facilitates better human-machine interaction and could be an efficient means to use intelligent robots in a more exible manner. In this paper, we report on our joint efforts at providing natural language access to the autonomous mobile two-arm robot Kamro. The robot is able to perform complex assembly tasks. To achieve autonomous behaviour, several camera systems are used for the perception of the environment during task execution. Since natural language utterances must be interpreted with respect to the robot's current environment, the processing must be based on a referential semantics that is perceptually anchored. Considering localization expressions, we demonstrate how, on the one

hand, verbal descriptions, and on the other hand, knowledge about the physical environment, i.e., visual and geometric information, can be connected to each other."
http://wwwipr.ira.uka.de/internal/detailed_publication.php?id=974907079

Thomas Längle, Tim C. Lüth, Eva Stopp, Gerd Herzog, and Gjertrud Kamstrup (1995). *KANTRA – A Natural Language Interface for Intelligent Robots.* International Conference on Intelligent Autonomous Systems, Karlsruhe, Germany, March. In Rembold et al. (eds.), *Intelligent Autonomous Systems*, IOS Press, pp. 357-364 (paper, 8 pages)

"The future use of advanced technical systems, for example robots in the area of service and maintenance, will lead to high demands for interaction between man-machine and machine-man to make these systems more easily accessible for human operators. On that account, natural language could be an efficient means to use robots in a more flexible manner. First, it is possible to convey information in a varying degree of condensation, on the other hand, communication can be performed on different levels of abstraction in an application- specific way. In order to fully exploit the capabilities of a natural language access, a dialogue-based approach will be used. In this article, we want to report about the joint efforts of the University of Karlsruhe and the University of the Saarland in providing natural language access to the autonomous mobile two-arm robot KAMRO which is being developed at IPR. A closer view at the description of the interface architecture will be given. It is documented how the integration of the man-machine interface into the control architecture of the robot should be performed in order to supply access to all internal information and models that are necessary for an autonomous behaviour."
http://wwwipr.ira.uka.de/internal/detailed_publication.php?id=975423485
http://wwwipr.ira.uka.de/internal/download.php?id=975423485&filetype=pdf