Computer Science cpsc590

Research Methods

in

Artificial Intelligence

Cristina Conati

Laboratory for Computational Intelligence

Overview

- ➡ What is AI? Different views.
- Brief LCI Overview
- Sample projects from Cristina's research group

What is Artificial Intelligence?

From: Russell S. and Norvig, P.``*Artificial Intelligence: A Modern Approach*." 2nd Ed., Prentice Hall, 2003

Systems that act like humans	**Systems that think rationally**
"The study of how to make computers do things at which, at the moment, people are better"(Rich and Knight, 1991)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985).
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning"(Bellman, 1978)	"The branch of computer science that is concerned with the automation of intelligent behavior (Luger and Stubblefield 1993)

Systems that *act like humans*

- Turing test (1950): Can a human interrogator tell whether (written) responses to her (written) questions come from a human or a machine?
  - Natural Language Processing
  - Knowledge Representation
  - Automated Reasoning
  - Machine Learning
- Total Turing Test (extended to include physical aspects of human behavior)
  - Computer Vision
  - Robotic

Has any AI System Passed the Tutoring Test?

- Not the full blown one (see http://www.loebner.net/Prizef/loebner-prize.html)
- Variations restricted to specific tasks requiring *some form* of intelligence
- Do we really want an intelligent system to act like a human?

ALICE

Generating Multimedia Presentations

Zhou, Wen, and Aggarwal. A Graph-Matching Approach to Dynamic Media Allocation in *Intelligent Multimedia Interfaces*. Best Paper Award at Intelligent User Interfaces 2005.

- Algorithm to effectively allocate text and graphics in multimedia presentations
- Empirical Validation
  - System (RIA) output on 50 user queries (real estate and tourist guide application)
  - Media allocation on same queries by two multimedia UI designers
  - Third expert "blindly" ranked all responses
- Results
  - RIA best/co-best in 17 cases
  - Minor differences in 28 of the remaining 33 cases

Why Replicate Human Behavior, Including its "Limitations"?

- AI and Entertainment
  - E.g. *Façade*, a one-act interactive drama
    http://www.quvu.net/interactivestory.net/#publications
- Sometime these limitations can be useful
  - E.g. Supporting Human Learning via Peer interaction
    (Goodman, B., Soller, A., Linton, F., and Gaimari, R. (1997) Encouraging Student Reflection and Articulation using a Learning Companion. *Proceedings of the AI-ED 97 World Conference on Artificial Intelligence in Education*, Kobe, Japan, 151-158.

What is Artificial Intelligence?

Systems that act like humans	**Systems that think rationally**
"The study of how to make computers do things at which, at the moment, people are better"(Rich and Knight, 1991)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985).
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning"(Bellman, 1978)	"The branch of computer science that is concerned with the automation of intelligent behavior (Luger and Stubblefield 1993)

Systems That *Think* Like Humans

- Use Computational Models to Understand the Actual Workings of Human Mind
  - Devise/Choose a sufficiently precise theory of the mind
  - Express it as a computer program
  - Check match between program and human behavior (actions and timing) on similar tasks
- Tight connections with *Cognitive Science*

Some Examples

- Newell and Simon's GPS (General Problem Solver, 1961) to test means-end approach as general problem solving strategy
- John Anderson's ACT-R cognitive architecture (http://act-r.psy.cmu.edu/)
  - Anderson, J. R. & Lebiere, C. (1998). *The atomic components of thought.* Erlbaum;
  - Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y . (2004). An integrated theory of the mind. *Psychological Review 111,* (4). 1036-1060.
- SOAR cognitive architecture (http://sitemaker.umich.edu/soar)
  - Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press.

ACT-R Models for Intelligent Tutoring

- One of ACT-R main assumptions:
  - Cognitive skills (procedural knowledge) are represented as production rules

Eq: 5x+3=30 ; Goals: [Solve for x]

- Rule: To solve for x when there is only one occurrence, unwrap (isolate) x.

Eq:5x+3=30 ; Goals: [Unwrap x]

- Rule: To unwrap ?V, find the outermost wrapper ?W of ?V and remove ?W

Eq: 5x+3=30; Goals: [Find wrapper ?W of x; Remove ?W]

- Rule: To find wrapper ?W of ?V, find the top level expression ?E on side of equation containing ?V, and set ?W to part of ?E that does not contain ?V

Eq: 5x+3=30; Goals: [Remove "+3"]

- Rule: To remove "+?E", subtract "+?E" from both sides

Eq: 5x+3=30; Goals: [Subtract "+3" from both sides]

- Rule: To subtract "+?E" from both sides ….

Eq: 5x+3-3=30-3

Model Tracing

- Given a rule-based representation of a target domain (e.g. algebra),
- an "expert model" can *trace* student performance by *firing* rules and do a stepwise comparison of rule outcome with student action
- Mismatches signal incorrect student knowledge that requires tutoring
- *Knowledge tracing* extends model tracing to assess probability that a student knows domain rules given observed actions
- These models showed good fit with student performance, indicating value of the ACT-R theory
- Also, the *Cognitive Tutors* based on this model are great examples of AI success – used in thousands of high schools in the USA

What is Artificial Intelligence?

Systems that act like humans	Systems that think rationally
"The study of how to make computers do things at which, at the moment, people are better"(Rich and Knight, 1991)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985).
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning"(Bellman, 1978)	"The branch of computer science that is concerned with the automation of intelligent behavior (Luger and Stubblefield 1993)

Systems that *Think Rationally*

- ***Logic***: formalize *right* thinking, i.e. irrefutable reasoning processes. Example:
  - Knowledge Base (KB)
    - E.g. *facts* (individuals) and *predicates* (relationships)
  - reasoning rules that can derive all correct conclusions from KB
    - E.g. *modus ponens:* `IF A -> B and A then B`
- Logistic tradition in AI aims to build computational frameworks based on logic.
- Then use these frameworks to build intelligent systems

Systems that *Think Rationally*

- Main Research Problems/Challenges
  - Proving *Soundness* and *Completeness* of various formalisms
  - How to represent often *informal* and *uncertain* domain knowledge and formalize it in logic notation
  - Computational Complexity

- Tradeoff between *expressiveness* and *tractability* in logic-based systems H. J. Levesque and R. J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. Computational Intelligence, 3(2):78--93, 1987.

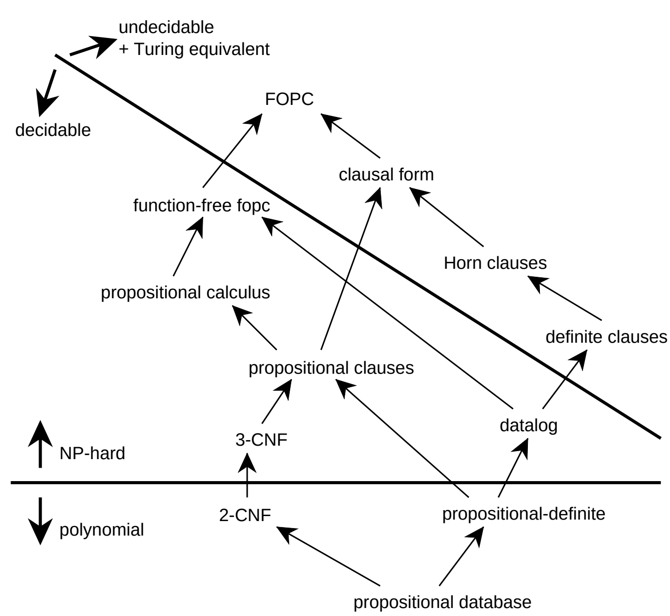

Current Research Areas

From topics list of IJCAI '05 (Major AI conference)

- abduction
  automated reasoning
  common-sense reasoning
  computational complexity
  description logics
  knowledge representation
  logic programming
  model-based reasoning
  nonmonotonic reasoning
- Ontologies
  qualitative reasoning
  reasoning about actions and change
  temporal reasoning
  theorem proving

What is Artificial Intelligence?

Systems that act like humans	**Systems that think rationally**
"The study of how to make computers do things at which, at the moment, people are better"(Rich and Knight, 1991)	"The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985).
Systems that think like humans	**Systems that act rationally**
"The automation of activities that we associate with human thinking, such as decision making, problem solving, learning"(Bellman, 1978)	"The branch of computer science that is concerned with the automation of intelligent behavior (Luger and Stubblefield 1993)

Systems that *Act Rationally*

- In the "*think rationally*" approach focus was on correct inference
- But more is needed for *rational behavior,* e.g.
  - How to behave when there is no provably correct thing to do (i.e. reasoning under uncertainty)
  - Fully reactive behavior (instinct vs. reason)

AI as Study and Design of Intelligent Agents

- An ***intelligent agent*** is such that
  - Its ***actions*** are ***appropriate*** for its goals and circumstances
  - It is ***flexible*** to changing environments and goals
  - It ***learns*** from experience
  - It makes ***appropriate choices*** given ***perceptual limitations*** and ***limited resources***
- This definition drops the constraint of *cognitive plausibility*
  - Same as building flying machines by understanding general principles of flying (aerodynamic) vs. by reproducing how birds fly
- Normative vs. Descriptive theories of Intelligent Behavior

Intelligent Agents in the World

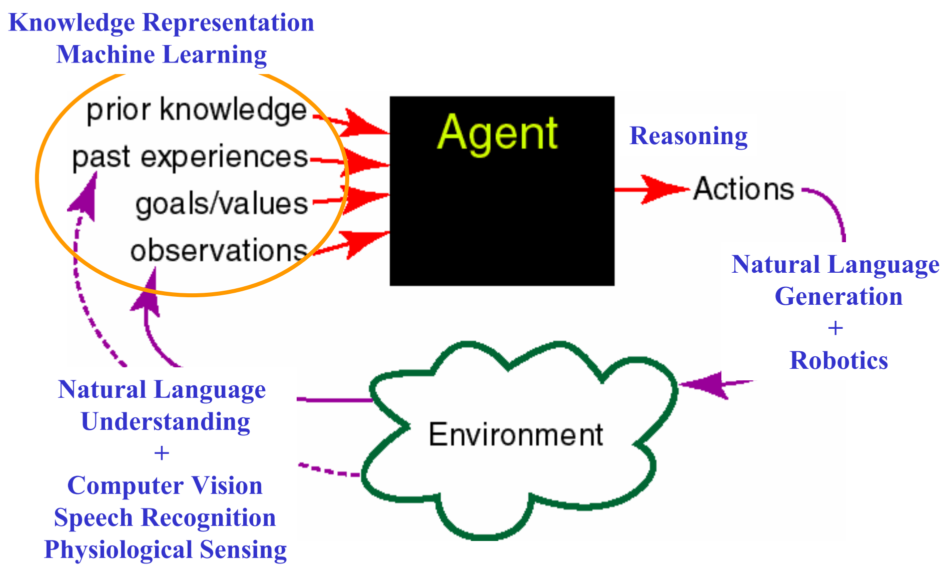

Approach

- ***Scientific goal***: to understand the principles that make intelligent behavior possible, in natural or artificial agents
- ***Engineering goal***: to specify methods for the design of useful, intelligent artifacts.
- ***Methodology***: design, build and experiment with computational systems that act intelligently.

Some Research Areas

- Game Theory (Billings, Burch, Davidson, Holte, Schaeer, Schauenberg, and Szafron. Approximating Game-Theoretic Optimal Strategies for Full-scale Poker, In Proc. IJCAI-03
- Reasoning Under Uncertainty (Judea Pearl, Probabbilistic Reasoning in Intelligent Systems, Networks of Plausible Inferences, Morgan Kaufmann)
- Decision Theory: rational decision making *under uncertainty* as finding actions or *policies* with maximum expected utility
  - Horvitz, Breese, and Henrion, Decision Theory in Expert Systems and Artificial Intelligence, J. of Approximate Reasoning, *Special Issue on Uncertainty in AI*, 2:247-302
- Statistical Approaches for
  - Machine Learning (Hastie, Tibshirani, and Friedman. The Elements of Statistical Learning: Data Mining, Inference and Prediction. Springer Series in Statistics, Springer-Verlag, 2001.)
  - Natural Language Processing (Manning and Schütze, *Foundations of Statistical Natural Language Processing*, MIT Press. Cambridge, MA: May 1999)

Trends

- AI is becoming more and more *interdisciplinary*
  - Philosophy
  - Mathematics
  - Economics
  - Neuroscience
  - Psychology
  - Computer Engineering
  - Linguistics

Trends

- AI is becoming more and more *rigorous*
  - Rely on existing theories instead of building new ones
  - Base claims on rigorous theorems or hard experimental evidence
  - Show relevance to real world applications rather than toy examples

Some Real World Applications

1999 Co-Winner of NASA's Software of the Year Award

It's one small step in the history of space flight. But it was one giant leap for computer-kind, with a state of the art artificial intelligence system being given primary command of a spacecraft. Known as Remote Agent, the software operated NASA's Deep Space 1 spacecraft and its futuristic ion engine during two experiments that started on Monday, May 17, 1999. For two days Remote Agent ran on the on-board computer of Deep Space 1, more than 60,000,000 miles (96,500,000 kilometers) from Earth. The tests were a step toward robotic explorers of the 21st century that are less costly, more capable and more independent from ground control.

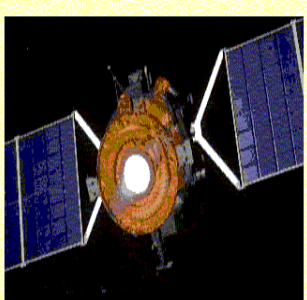

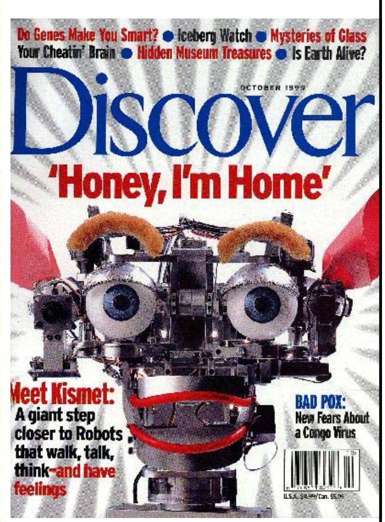

Some Real World Applications

- ALVINN: learns to steer a vehicle on single lane road
  - Has driven up to 2850 miles on public highways, as well on dirt roads and bike paths, being in control for 98% of the time
- Deep Blue: First program to defeated chess world champion
- HipNav:
  - computer vision techniques to create 3D model of the patient internal anatomy
  - robotics control to guide the insertion of hip replacement prosthesis
- Cognitive Tutors

Overview

- What is AI? Different views.
- ➡ Brief LCI Overview
- Sample projects from Cristina's research group

Laboratory for Computational Intelligence at UBC

(http://www.cs.ubc.ca/nest/lci/)

- 14 faculty
- 10 associate faculty
- 4 postdocs
- 26 Graduate Students

Research Areas

- Computer Vision (Little, Lowe, Murphy, Woodam)
- Constraints-Based Reasoning (Mackworth)
- Intelligent User Interfaces (Carenini, Conati)
- Machine Learning (de Freitas, Murphy)
- Multi-Agents Systems (Leighton-Brown, Mackworth)
- Reasoning Under Uncertainty (Conati, Murphy Poole)
- Robotics (Little, Mackworth, MacLean)

Why “Computational” Intelligence?

- “Artificial” can be interpreted as ***fake*** or ***simulated***

- the goal is to understand ***real*** intelligent systems by synthesizing them.