

Research Methods In Artificial Intelligence

Alan Mackworth

Laboratory for Computational Intelligence
Computer Science, UBC
<http://www.cs.ubc.ca/~mack>
mack@cs.ubc.ca

October 6, 2003

Talk Outline

- What is AI?
- What is I?
- AI: Science or Engineering or Marketing?
- AI Schools of Thought
- What's an AI result?
- The Deadly Sins of AI Research
- Doing the Work
- Getting the Word Out (and In)
- The Bottom Line

Debate Topic

Motion: Resolved that all research in AI is contaminated by at least one of the Deadly Sins.

What is AI?

What is Computational/Artificial Intelligence? The study of the design of intelligent agents.

An agent is something that acts in an environment.

An intelligent agent is an agent that acts intelligently. (Duh!)

AI is task not discipline-focussed: we scavenge for theories in many different disciplines: CS, psychology, control theory, applied math, probability, economics, linguistics, philosophy,

That's what makes it so much fun!

BUT it makes evaluation harder since individual theories usually have unique evaluation criteria.

What is I?

An intelligent agent is:

Proactive An agent achieves goals, implicit or explicit. Its behaviour is teleological, planned and future-oriented.

Reactive An agent perceives and reacts to environmental change. Its behaviour is causal and past-determined.

Model-based It uses models of the world to guide its perception and behaviour.

Learning-oriented It acquires new behaviours and new models.

Rational It reasons, solves problems and uses tools.

Social It collaborates, cooperates, commits and competes with other agents.

Linguistic It communicates and coordinates using language.

Situated It is embedded or situated in a world to which it is coupled. It is particular not universal.

Constraint-based It satisfies and optimizes multiple external and internal constraints.

A researcher's approach to AI research will be determined by how (s)he ranks the relative importance of these criteria.

What is LCI?

LCI = The UBC Laboratory for Computational Intelligence
<http://www.cs.ubc.ca/labs/lci>

AI research in LCI presented tomorrow in the Grad Seminar
October 7, 2003 at 12:30 in MCLD 214

AI: Science or Engineering or Marketing?

- Scientific goal: to understand the computational principles that make intelligent behavior possible, in natural or artificial systems.
- Engineering goal: to specify methods for the design of useful, intelligent artifacts.
- Marketing goal: to make money by selling products or ideas

Analogy between studying flying machines and thinking machines.

Analysis or design? Science requires analytic theory and experimental design. We build prototypes for experiments but they are not product designs.

Some AI Application Areas

- Perception - Vision
- Reasoning & Planning - Logic and Probability
- Robotics
- Language understanding and generation - Speech & text
- Multi-agent systems
- Affective systems
- Games: Checkers, Chess, Poker, Soccer, Video games, ...
- User Interfaces - Education
- Cognitive Science

Some AI Schools of Thought

- Cybernetics
- Pattern Recognition
- 'GOFAIR'
- 'Insect AI'
- Behaviour-based
- Neural Nets
- A-life
- Game theoretic
- Agent-oriented

“When you have a hammer, everything looks like a nail.”

Some Current Trends: Situated Agents

- Coupled to the environment
- Dynamic environment: other agents – social roles
- Tight coupling of perception and action
- Embody environmental constraints by design, evolution, learning
- Situated communication & perception (external & internal)
- On-line architectures: transductions
- Reason about agents (correctness) as well as in agents
- Unitary model: “No hybrid models for hybrid systems”

What's an AI result?

Here are the two Distinguished Paper Award winners at the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03):

Mark A. Paskin. Thin Junction Tree Filters for Simultaneous Localization and Mapping. In Proc. IJCAI-03, pp. 1157-1164, 2003.
<http://www.cs.berkeley.edu/~paskin/>

Darse Billings, Neil Burch, Aaron Davidson, Rob Holte, Jonathan Schaeffer, Terence Schauenberg, and Duane Szafron. Approximating Game-Theoretic Optimal Strategies for Full-scale Poker, In Proc. IJCAI-03, pp. 661-668, 2003.
<http://www.cs.ualberta.ca/~jonathan/Papers/ai.2003.html>

AI and Approximation

These two papers study very different tasks and problems but they share a focus on approximate reasoning. This is not an accident. Much AI deals with NP-hard problems so approximate solutions are often required. These papers study approximate solutions from both theoretical and practical perspectives. Often AI deals with inverse problems (which are intrinsically hard) not direct problems. Compare graphics and vision. So approximation is required for scaling.

The Nine Deadly Sins of AI Research

1. Look ma, no hands
2. Plato (and Otalp)
3. Little boxes
4. Gee, what shiny bricks
5. One day my son
6. Got no time for you right now
7. Tablets of stone
8. Rally round the flag boys
9. It's all really only

Old Sins and New Confessions, Patrick Hayes and Kenneth Ford, AI Magazine

20(2): Summer 1999, 128

<http://www.aaai.org/Library/Magazine/Vol20/20-02/vol20-02.html>

Assessing the State of the Art

Read critically.

Journals: AIJ, JAIR, IJCV, IJRR, PAMI, Constraints, CI,

Conferences: IJCAI, AAI, ICCV, ICRA, ECAI, UAI,

Talk to people and listen! Go to workshops and conferences.

Network.

Doing the Work

Research questions: What's new? What's different? Why is it important? Demonstrate results practically and/or theoretically.

How to do Research At the MIT AI Lab

by: a whole bunch of current, former, and honorary MIT AI Lab graduate students

David Chapman, Editor September, 1988.

<http://www.cs.indiana.edu/mit.research.how.to/mit.research.how.to.html>

The biggest challenge is choosing the right task - just hard enough but no harder - that focusses on stretching current theory, e.g. soccer.

Getting the Word Out (and In)

Telling the world is integral to the research enterprise.

Pat Winston's heuristics for giving talks:

<http://www.cs.dartmouth.edu/~brd/Teaching/Giving-a-talk/phw.html>

And listening to the feedback is too!

Insert standard advice here on writing CS theses, conference and journal papers.

Many audiences for AI research. Target the right audience with the right message. Go to workshops and conferences. Give posters and talks. Network.

The Bottom Line

- AI is an exciting area but you have to dig behind the hype.
- Keep the science/engineering distinction clear.
- Given its multidisciplinary nature you have to keep acquiring new tools.
- Understand the development of the field and choose your approach carefully.
- Build on the most exciting work.
- Get the word out there.
- Have fun: that's the objective function!