Cognitive Architecture for Robots[1]

OR65: AI and ML Theme, 2023

Mohan Sridharan

Reader in Cognitive Robot Systems
Intelligent Robotics Laboratory
School of Computer Science, University of Birmingham, UK
m.sridharan@bham.ac.uk

September 14, 2023

[1] Hasra Dodampegama (Univ. of Birmingham, UK); Michael Gelfond (Texas Tech, USA); Ben Meadows, Rocio Gomez, Tiago Mota, Heather Riley (Univ. of Auckland, NZ); Saif Sidhik, Jeremy Wyatt (Univ. of Birmingham, UK); Shiqi Zhang (SUNY Binghamton, USA), US ONR Awards N00014-13-1-0766, N00014-17-1-2434, N00014-20-1-2390; US AFOSR Award FA2386-16-1-4071; UK EPSRC Award EP/S032487/1; Honda Research Institute EU.

Research Objectives

Develop architectures and algorithms that:

- represent, reason, and act with different descriptions of domain knowledge and uncertainty.
  "Books are usually in the library"
  "I am 90% certain the robotics book is in the library"

- learn interactively and cumulatively from stream of noisy multimodal sensor cues.
  Learn actions, action capabilities, domain dynamics
  "Robot with weak arm cannot lift heavy box"

- enable designers to understand the behavior, and establish that it satisfies desired properties.
  Explainable agency (goals, beliefs), safety
  "What would happen if I dropped the the spoon on the table?"

Agents: Software, Robots, and Humans!

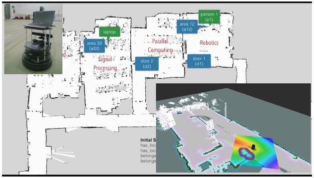

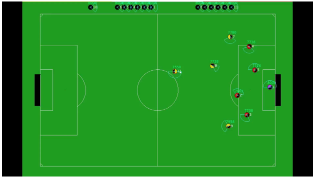

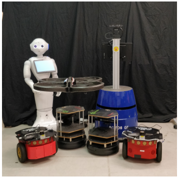

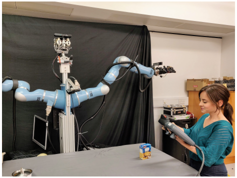

Collaboration without Prior Coordination

Example Example

- Limited prior knowledge of other agents/robots; observable state but no (limited) communication.
- **State of the art**: data driven methods.
  - Probabilistic and/or deep network-based models.
  - Estimate behavior of other agent "types", optimize actions using long experience history. Transparency, accessibility issues?

Reuth Mirsky, Ignacio Carlucho, Arrasy Rahman, Elliot Fosong, William Macke, Mohan Sridharan, Peter Stone, Stefano V Albrecht. **A Survey of Ad Hoc Teamwork: Definitions, Methods, and Open Problems**, arXiv:2202.10450, 2022.

Core Ideas and Claims

- Cognitive systems inspired by human cognition and control.
- Represent, reason, act, learn jointly: different abstractions and schemes; formal coupling of logician, statistician, explorer.
- Knowledge-based and data-driven reasoning, control, learning.
- "Here and there" non-monotonic reasoning and learning.
- Ecological rationality: adaptive satisficing, algorithmic model of heuristic methods, anticipatory models.
- Explainable agency: relational descriptions that make contact with human concepts such as beliefs and goals.

Architecture Overview

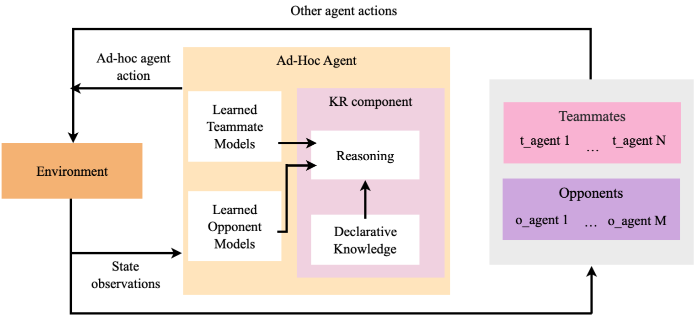

Exploit complementary strengths of non-monotonic logical reasoning, probabilistic reasoning, and interactive learning.

Logician: Coarse-resolution Reasoning

- Abstract action language description ($\mathcal{D}_C$) of transition diagram.
- Statics: $next\_to(region, region)$; Fluents: $loc(agent, region)$; Actions: $move(agent, place)$, $shoot(agent, agent)$.
- Axioms (causal law, state constraint, executability condition):

$$move(ag_1, R) \ \textbf{causes} \ loc(ag_1) = R$$
$$\neg face(ag_1, d_2) \ \textbf{if} \ face(ag_1, d_1)$$
$$\textbf{impossible} \ shoot(ag_1, Ag) \ \textbf{if} \ not \ in\_range(ag_1, Ag)$$

- History ($\mathcal{H}_C$) of observations, actions, initial state defaults.
- Planning, diagnostics: Answer Set Prolog program $\Pi(\mathcal{D}_C, \mathcal{H}_C)$.
- Nonmonotonic logic: default negation, epistemic disjunction.

`¬ l`	l is believed to be false	`not l`	it is not believed that l is true
`p ∨ ¬ p`	is tautological	`p or ¬ p`	is not tautological

Statistician: Refine + Zoom + Execute

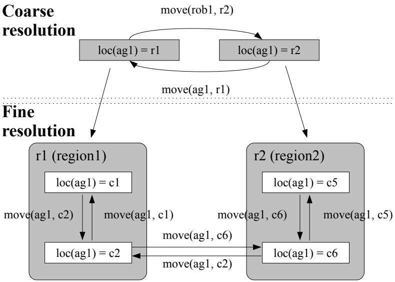

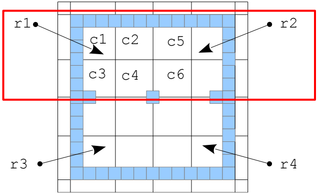

- Refinement: view domain description at finer resolution.
- Randomize and zoom to relevant part; formal guarantees.
- Probabilistic sequential decision making if needed: observations and outcomes update history.

Mohan Sridharan, Michael Gelfond, Shiqi Zhang and Jeremy Wyatt. **REBA: Refinement-based Architecture for Knowledge Representation and Reasoning in Robotics**. In Journal of Artificial Intelligence Research, 65:87-180, May 2019.

Explorer: Predictive Behavior Models

- Ecological rationality: adaptive satisficing (Herb Simon); behavior a function of cognition and environment.
- Algorithmic model of heuristics: ignore some information; make decisions more quickly, frugally, accurately.
- Identify attributes. Learn sparse models; match methods to domain characteristics: ensemble of fast and frugal trees.

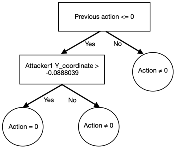

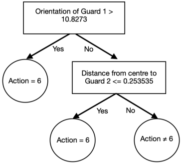

G. Gigerenzer. **What is Bounded Rationality?** Routledge Handbook of Bounded Rationality, 2020.
K. Katsikopoulos, O. Simsek, M. Buckmann, G. Gigerenzer. **Classification in the Wild: The Science and Art of Transparent Decision Making**, MIT Press, 2021.

Experimental Setup and Findings

- Train with hand-coded policies, test on deep network policies.
- Adaptation to different teammate and opponent types, different team sizes, different domains.
- Consider partial observability and limited communication.
- Learn predictive models from orders of magnitude fewer examples (10000 vs. 1M).
- Better performance (episode duration, wins, shooting accuracy) than state of the art data-driven systems.
- Relational descriptions of decisions and beliefs.

HFO: Prediction Accuracy + Goal Scoring

Agent Type	**Accuracy**
Helios	86.0%
Gliders	66.4%
Cyrus	77.6%
Aut	67.7%
Axiom	73.6%
Agent2D	71.9%

Version	**KAT** (%)	**PPAS** (%)	**PLAS** (%)
Limited (2v2)	79	80	80
Full (4v5)	30	20	20

Execution Video 1: Fort Attack

FA-result

Hasra Dodampegama and Mohan Sridharan. **Toward a Hybrid Architecture for Ad Hoc Teamwork**. In the AAAI Conference on Artificial Intelligence (AAAI), Washington DC, USA, February 7-14, 2023.

Execution Video 2: Half Field Offense

HFO-result

Hasra Dodampegama and Mohan Sridharan. **Knowledge-based Reasoning and Learning under Partial Observability in Ad Hoc Teamwork**. In Theory and Practice of Logic Programming, 2023.

Execution Example: Explanation

Human: ”Why did you move to (3,14) in step 0?”

Agent: ”Because attacker1 was not in range and I had to move to (4,14)”

Human: ”Why did you not move to (4,13) in step 2?”

Agent: ”Because attacker1 would not be in range”

Human: “What would have happened if you stayed in place in step 2?”

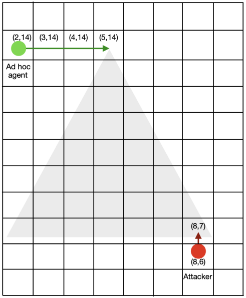

Long-term goal: eliminate attacker1 and protect fort.
Short term goal: getting to location (5, 14).

Hasra Dodampegama and Mohan Sridharan. **Collaborate and Explain on the Fly: Nonmonotonic Logical Reasoning and Incremental Learning for Ad Hoc Teamwork**. In International Workshop on Nonmonotonic Reasoning, 2023.

Contributions

- Step-wise refinement: simplifies design; confidence and scalability.
- Non-monotonic logical reasoning, probabilistic reasoning, and interactive learning inform and guide each other.
- Ecological rationality for reliable and efficient reasoning and learning.
- Interactive explanations of decisions constructed on demand.
- Explore scalability under constrained observation and communication; learn and revise knowledge.
- Explore explanation and composition of learned models and behaviors.

That's all folks!

PhD and **Research Fellow** positions available.

Extra: Changing-Contact Manipulation

Same principles applied to robot manipulation problem.

Changing contact manipulation

Saif Sidhik. An Online Framework for Changing-Contact Robot Manipulation, Doctoral Dissertation, School of Computer Science, University of Birmingham, UK, 2022.

Saif Sidhik, Mohan Sridharan, and Dirk Ruiken. **Towards a Framework for Changing-Contact Robot Manipulation**. In the International Conference on Intelligent Robots and Systems (IROS), Prague, Czech Republic, 2021.

Extra: Reasoning and Learning

Same principles used to combine knowledge-based and data-driven methods for reasoning and learning.

SAM Example

N. Gireesh, A. Agrawal, A. Datta, S. Banerjee, M. Sridharan, B. Bhowmick, and M. Krishna. **Sequence-Agnostic Multi-Object Navigation**. In the IEEE International Conference on Robotics and Automation (ICRA), London, UK, May 29-June 3, 2023.

A. Agrawal, R. Arora, A. Datta, S. Banerjee, B. Bhowmick, K.M. Jatavallabhula, M. Sridharan, and M. Krishna. **CLIPGraphs: Multimodal Graph Networks to Infer Object-Room Affinities**. In the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), Busan, Korea, August 28-31, 2023.

Heuristics Methods for Predictive Behavior Models

- Ecological rationality: adaptive satisficing; behavior a function of cognition and environment.
- Herb Simon's notion of bounded rationality.
- Adapted to focus on optimality (finance); heuristics and biases (psychology).
- Algorithmic model of heuristics: ignore some information; make decisions more quickly, frugally, accurately.
- Identify good attributes; learn predictive models. Match methods to domain characteristics.

Heuristics Methods for Predictive Behavior Models

- Identify good attributes; minimal, abstract, representation.

Description of attribute	**Number**
x position of agent	6
y position of agent	6
distance from agent to center of field	6
agents' polar angle with center of field	6
orientation of the agent	6
distance from agent to fort	6
distance to nearest attacker from fort	1
number of attackers not alive	1
previous action of the agent	1

Heuristics Methods for Predictive Behavior Models

- Match classes of heuristic methods to domain characteristics.

- Ensemble of fast and frugal trees.

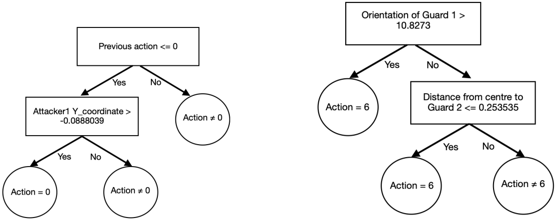

- Can learn and revise rapidly.

- Predict action choices of other agents.

Partial Observability and Communication

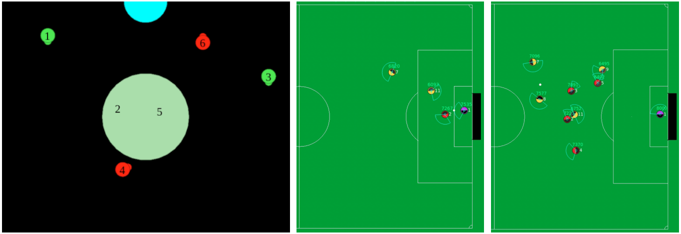

- Extend to consider partial observability.
- Communicate when, what, and to whom?
- Extend to more complex domains.

FA: Prediction Accuracy + Shooting Accuracy

Agent Model	Accuracy
Guard type 1	85.5%
Guard type 2	60.0%
Attacker type 1	86.9%
Attacker type 2	85.2%

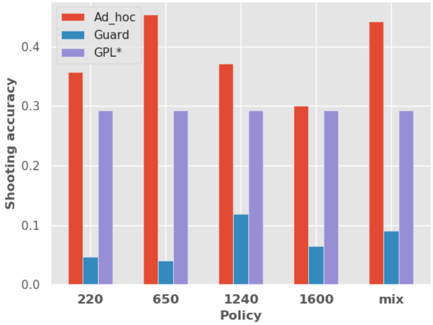

HFO: Prediction Accuracy + Goal Scoring

Agent Type	**Accuracy**
Helios	86.0%
Gliders	66.4%
Cyrus	77.6%
Aut	67.7%
Axiom	73.6%
Agent2D	71.9%

Version	**KAT** (%)	**PPAS** (%)	**PLAS** (%)
Limited (2v2)	79	80	80
Full (4v5)	30	20	20