Reflective, Deliberative Information Gathering (Doctoral Consortium)

Adam Eck
Department of Computer Science and Engineering
University of Nebraska-Lincoln
aeck@cse.unl.edu

ABSTRACT
In this research, we focus on reflective and deliberative solutions to improve information gathering by agents in complex environments to handle challenging properties such as uncertainty, partial observability, non-stationarity, and limited resources. We describe the core problems addressed by this research, our ongoing contributions towards solving these problems, and describe how we are applying our research to two real-world applications: personal assistants and intelligent survey systems.

Categories and Subject Descriptors
I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

Keywords
Information Gathering; Metareasoning; Observer Effect; Sharing

1. INTRODUCTION
Many real-world applications of computer systems benefit from the use of artificial intelligence and multiagent systems. For example, intelligent agents have found wide-ranging uses from personal assistants in educational systems [10] to robotics for area surveillance [13]. Here, agents are beneficial as they enable a system to achieve valuable properties such as reliability, scalability, robustness, consistency, efficiency, and effectiveness. To provide these benefits, intelligent agents are responsible for making autonomous, intelligent decisions in order to accomplish goals and complete tasks. However, the ability of agents to make decisions depends on the quality of information gathered by the agent from its environment through *sensing*: without good information, a rational agent could make wrong decisions and thus fail to accomplish its goals and complete its tasks. Such sensing is especially difficult in complex, real-world environments due to challenging environment properties such as uncertainty, partial observability, non-stationarity, and limited resources.

To improve agent reasoning, this research focuses on **reflective, deliberative information gathering** by intelligent agents. An agent behaves so by carefully considering its current knowledge, the knowledge required of its decisions, and the state of its environment in order to know how, when, and where to sense so that it gathers appropriate information. By being ***reflective***, an agent can *self-evaluate* its informational needs and performance in order to *adapt* and *learn* as it faces new decisions in an uncertain environment. By being ***deliberative***, an agent can *selectively* choose how to gather information and *intentionally* optimize its information gathering and thereby make the best decisions possible, as opposed to considering sensing as a *secondary* behavior only reactively performed as necessary or as a by-product of other activities, which would potentially lead to suboptimal information gathering in complex environments. This research (1) extends classical metareasoning from decisions about reasoning control to decisions about sensing control to further support reasoning, and (2) extends active sensing/perception [14] to be more introspective about agent performance and needs in order to encourage improved adaptation over time.

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

2. CORE PROBLEMS
To support reflective, deliberative information gathering, this research addresses three core problems:

Analysis Problem: How should an agent *reflect upon the benefits and costs* of performing various sensing actions with respect to gathering information?

We have explored how to embed different types of rewards and guides within an agent's reasoning in order to allow the agent to reflect on and analyze the benefits and costs of performing information gathering actions, then deliberatively choose sensing actions to balance those benefits and costs and improve reasoning. This research has produced decision-theoretic solutions using probabilistic planning for uncertain environments [1, 2, 4, 7] that optimize sensing by considering (i) the value of information an agent might gather with respect to the current decisions the agent must make against (ii) the costs of gathering such information. For instance, we've developed a novel context-aware heuristic search algorithm for POMDPs [9] that improves planning and subsequent agent goal achievement in complex, highly uncertain domains [1]. This algorithm intelligently chooses between actions that either rapidly reduce the agent's uncertainty or exploit agent knowledge to maximize rewards and minimize costs, depending on the agent's most pressing needs. In this manner, the agent finds usable plans an order of magnitude faster than previous state-of-the-art heuristic search algorithms in the types of highly uncertain environments agents are likely to experience in real-world applications. We've also formulated an extension of potential-based reward shaping (PBRS) [11] within POMDPs to provide immediate metareasoning (e.g., assessing certainty in agent beliefs) to achieve greater long-term rewards (and thus more optimal actions) within the quick, short-term planning required of agents operating in real-world environments on computationally limited devices, such as embedded sensors and robotics [7].

Environment Impact Problem: How can an agent *mitigate the changes to its environment* caused by sensing that have *lasting impacts* on both the information gathered and the ability of the agent to accomplish its tasks?

We have also addressed an important consequence of agent sensing when *stateful resources* are required for information gathering: the **Observer Effect of agent sensing** [2]. Specifically, as an agent interacts with a stateful resource (e.g., a human user with variable frustration or cognitive loads), the agent can change the state of the resource by performing sensing actions (e.g., further frustrating a user through interruptions), which causes dynamic (rather than fixed) costs to the agent based on the state of the resource. More importantly, the quality of information gathered by a stateful resource depends on its current state -- a frustrated user might hurriedly respond with less accurate or fewer information to quickly return to her original task when interrupted. Therefore, agents must be mindful of the internal state of resources used during sensing in order to gather the best quality information at the lowest cost. Reflective awareness of these impacts and deliberative choices to sense appropriately enable us to build better agents capable of more efficient and effective interactions with human users, while also improving the end-user experience.

Information Sharing Problem: How can agents leverage multiagent cooperation in order to *share information when information gathering is limited* (e.g., agents have limited sensors or resources)?

Finally, we have investigated the dynamics of information flow through multiple cooperative agents working together as they share information, focusing thus far on challenging domains with localized phenomenon observed by only a small subset of the agents within large cooperative teams (e.g., observing individual users of a large mixed initiative software system), requiring large team information sharing [8, 12] to reach consistent, accurate shared beliefs. Our research has produced solutions to overcome a challenging problem: the **institutional memory problem** where large portions of the team of agents become stuck with outdated beliefs as the non-stationary environment changes (e.g., changing user preferences or goals) [3]. We have proposed two algorithms for mitigating this problem: (1) a change detection and response algorithm where agents work together within local sub-teams to quickly detect changes to the observed phenomenon, and (2) a forgetting-based algorithm, where agents independently use belief decay to maintain up-to-date beliefs to avoid problems caused by faulty agents or malicious information. Both solutions successfully avoid the institutional memory problem and lead to consistent, accurate beliefs through the team as the environment changes. We are currently working to extend this work to additional models of knowledge, communication, and environment characteristics.

3. APPLICATIONS

Altogether, this research is currently being applied to two real-world applications within distinct, interdisciplinary domains: (1) **personal assistants** [5] for **computer supported, collaborative learning** [10], where human users interact with an agent-powered intelligent user interface to gather information and produce collaborative writing assignments in a wiki-based software system [6], and (2) **survey informatics** and **adaptive survey instruments**, where intelligent agents process information gathered from human users and adapt the surveys used to interact with human sources of information in order to reduce problems such as users *satisficing* by providing suboptimal information (e.g., repeatedly picking the first response to finish faster) or *breaking off* and failing to complete the survey (reducing information collection). These applications provide real world, complex environments for studying the three problems identified above, and solutions to these problems produced by our research advance both the ***state-of-the-art in multiagent information gathering*** research and ***improve human-agent interactions*** in computer systems, as well as produce systems deployed for use in the real world.

Beyond these two applications in human-computer (and human-agent) interactions, this research has potential broader impacts in domains such as robotics (e.g., search and rescue robots exploring unknown spaces), wireless sensor networks (e.g., limited network bandwidth monitoring and control), and cyber-physical systems (e.g., controlling limited numbers of sensors to form consistent beliefs across many devices).

4. ACKNOWLEDGMENTS

This research has been supported by a NSF Graduate Research Fellowship (DGE-054850), a U.S. Department of Education GAANN Fellowship (P200A040150), as well as grants from the NSF (DBI-0743783, DEB-0953766, SBES-1228937), and has been partially completed utilizing the Holland Computing Center of the University of Nebraska.

5. REFERENCES

[1] Eck, A. and Soh, L.-K. 2014. Online Heuristic Planning for Highly Uncertain Domains. *Proc. of AAMAS'14*.

[2] Eck, A. and Soh, L.-K. 2013. Observer Effect from Stateful Resources in Agent Sensing, *Autonomous Agents and Multi-agent Systems*, 26(2), 202-244.

[3] Eck, A. and Soh, L.-K. 2013. Dynamic Facts within Large Team Information Sharing (Extended Abstract). *Proc. of AAMAS'13*.

[4] Eck, A. and Soh, L.-K. 2012. Evaluating POMDP Rewards for Active Perception (Extended Abstract). *Proc. of AAMAS'12*.

[5] Eck, A. and Soh, L.-K. 2012. Adaptive Knowledge Assistants. *Proc. of HAIDM'12*.

[6] Eck, A., Soh, L.-K., and Brassil, C. 2013. Supporting Active Wiki-based Collaboration. *Proc. of CSCL'13*, 176-183.

[7] Eck, A., Soh, L.-K., Devlin, S., and Kudenko, D. 2013. Potential-Based Reward Shaping for POMDPs (Extended Abstract). *Proc. of AAMAS'13*.

[8] Glinton, R., Scerri, P., and Sycara, K. 2010. Exploiting scale invariant dynamics for efficient information propagation in large teams. *Proc. of AAMAS'10*, 21-28.

[9] Kaelbling, L.P., Littman, M.L., and Cassandra, A.R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101, 99-134.

[10] Khandaker, N., Soh, L.-K., Miller, L. D., Eck, A., & Jiang, H. 2011. Lessons Learned from Comprehensive Deployments of Multiagent CSCL Applications I-MINDS and ClassroomWiki. *IEEE Transactions on Learning Technologies*, 4(1), 47-58.

[11] Ng, A.Y., Harada, D., and Russell, S. 1999. Policy invariance under reward transformations: Theory and application to reward shaping. *Proc. of ICML'99*, 278-287.

[12] Pryymak, O., Rogers, A., and Jennings, N.R. 2012. Efficient opinion sharing in large decentralized teams. *Proc. of AAMAS'12*.

[13] Spaan, M.T.J., Veiga, T.S., and Lima, P.U. 2010. Active cooperative perception in networked robotic systems using POMDPs. *Proc. of IROS'10*, 4800-4805.

[14] Weyns, D., Steegmans, E., and Holvoet, T. 2004. Towards active perception in situated multi-agent systems. *Applied Artificial Intelligence*, 18, 867-883.