Information Sharing for Care Coordination (Doctoral Consortium)

Ofra Amir
School of Engineering and Applied Sciences, Harvard University, USA
oamir@seas.harvard.edu

ABSTRACT

Teamwork and care coordination are of increasing importance to health care delivery and patient safety and health. This thesis aims at developing agents that are able to make intelligent information sharing decisions to support a diverse, evolving team of care providers in constructing and maintaining a shared plan that operates in uncertain environments and over a long time horizon.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*

Keywords

Multi-Agent Planning; Information Sharing; Healthcare

1. INTRODUCTION

The health care literature argues compellingly that teamwork is of increasing importance to health care delivery, and improved care coordination is essential to improving patient safety and health. The lack of effective mechanisms to support health care providers in coordinating care is a major deficiency of current health care systems [9]. This thesis is part of a broader project that aims to develop intelligent, autonomous multi-agent systems that work as a team supporting a diverse, evolving team of providers caring for children with complex conditions [1]. The agents will support providers in formulating a shared "care plan" that operates on multiple time scales and in uncertain environments, deploying that plan in their delivery of care, and monitoring and revising it as needed.

The care team for a child with a complex condition is diverse and broad in scope, including not only physicians but also other types of care providers (e.g., therapists) and others who work with the child in various settings (e.g., teachers). The group of providers may change over the years, whether as a result of personnel changes or because the child's condition or developmental stage raise different needs. Thus, caregivers' involvement with the child may be continuous or intermittent, long or short term. Providers differ in their expertise and address different aspects of a child's condition. Furthermore, the care team for these children seldom all come together to discuss the child's care. While these factors make plan coordination and management very complex, it is crucial to the quality of care that this group acts as a team. To do so requires effective mechanisms for information sharing between team members. The focus of this thesis is developing supporting agents that are capable of making intelligent information sharing decisions.

[1]This project accords with the vision described in the paper "Collaborative Health Care Plan Support" in the challenges and visions track of the technical conference.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

2. INFORMATION SHARING

Our research on information sharing for care team coordination aims to develop agents capable of assisting care providers in ensuring that their individual treatments plans mesh and that those actions directed at short-term goals are compatible with longer-term goals. Insufficient communication among team members can lead to missing required treatment actions or conflicting actions. Agents could support caregivers with multiple responsibilities, and therefore limited time, by identifying from the large, heterogeneous body of information each has individually that portion most relevant to share, and determining those caregivers with whom it is important to share it.

The heterogeneity of the care team for a child with a chronic condition—in expertise, roles, knowledge about a child's condition and concern with the child's care—raises many challenges. Agents supporting such care teams will themselves operate in different contexts and have different beliefs about the state of the world and what each other and the people whom they support are doing. They will need strategies for resolving differences among contexts and for learning other agents' contexts, including care providers' intended actions and beliefs.

Prior multi-agent systems efforts e.g., Electric Elves [3], CALO [14], RADAR [5] have developed multi-agent planning technologies for personal assistant agents; however they focused on the support of a single individual. Some prior work, including Coordinators [13], has addressed collaboration among personal assistant agents, but there are key differences between these efforts and the goal of agents supporting care providers. First, in prior work, the collaborating agents belonged to a single organization, whereas collaborators in the health care domain usually are not. Second, prior work has considered agents that, though distributed, are tightly coupled in their efforts. In the health care do-

main, different providers operate semi-independently and have many competing tasks, including caring for other patients. Their plans are only loosely coupled, but doing the right thing when they interact is essential to plan success.

Existing techniques for collaborative multi-agent planning are not fully able to handle essential characteristics of the care plan setting and often do not take into account communication mechanisms and costs. Classical planning and agent frameworks such as BDI agents [10] assume a closed world where the operators and goals are defined and fixed from the start. Dec-POMDP models [2] address uncertainties about action outcomes and about states, but are intractable for long horizon plans and are not suited to incorporate new actions and agents as the care plan evolves, as they assume that a complete model of states and transitions is given in advance. Theories of teamwork and collaboration [6, inter alia] support collaborative multi-agent planning, but assume a fixed action library and set of agents. A range of recent work on information exchange and communication algorithms for multi-agent settings has defined models for communication within teamwork [8, 11, 4], but these models have been implemented and evaluated only in environments much simpler than the care coordination domain. They also typically tightly limit communications options and make modeling assumptions, including constant team membership, which do not hold in this domain.

Our planned approach will build on the work described above, as well as that on modeling collaboration [6], helpful behavior [7] and interruption management [8, 12] It will proceed in stages from simpler to more complex settings to address the information-sharing challenges of the care coordination domain. We are currently evaluating information sharing agents in a domain introduced by Roth et al. [11] in which players need to infer whether they are in Colorado or Wyoming based on terrain observations and meet in an agreed location. Although this domain is much simpler than the care coordination setting, it shares some of its characteristics, including costly communication, agents who differ in their observations and uncertainty regarding the state they are in. Prior work has used this domain to evaluate computer agents interacting with each other [11] and a computer agent interacting with a human teammate [4]. In our setting computer agents will be designed to support human players by making information sharing decisions. This game setting is analogous to the healthcare setting in which care providers ultimately make treatment decisions, but could benefit from a supporting computer agent that would reduce the burden of communication by identifying the relevant information for each provider and sharing this information at the right time.

In the next stage, we will design and test agents in a more complex game which has more of the characteristics of the care coordination scenario. To this end we have designed a game in the ColoredTrails framework [1]. In this game human players will need to agree on a set of shared non-conflicting goals to pursue on a grid board (corresponding to choosing between alternative treatment options in the healthcare domain). In the game there is uncertainty about the utility of each goal, and agents obtain more accurate information as they progress towards the goal (similarly to physicians who can only see whether a treatment works as they start pursuing it). Moreover, players cannot observe the positions and actions of other agents (as care providers are typically unaware of all actions performed by other providers). This abstract environment will enable us to test different representations and agent designs for supporting information sharing in various settings. Finally, we will incorporate the decision-making mechanisms in agents designed to operate in the care coordination domain and test their abilities to support caregivers.

3. REFERENCES

[1] The colored trails project. `http://www.eecs.harvard.edu/ai/ct`.

[2] D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, pages 819–840, 2002.

[3] H. Chalupsky, Y. Gil, C. Knoblock, K. Lerman, J. Oh, P. D.V., T. Russ, and M. Tambe. Electric elves: Agent technology for supporting human organizations. *AI Magazine*, 23(2):11–24, 2002.

[4] A. Frieder, R. Lin, and S. Kraus. Agent-human coordination with communication costs under uncertainty. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pages 1281–1282, 2012.

[5] D. Garlan and B. Schmerl. The RADAR architecture for personal cognitive assistance. *Intrnational Journal of Software Engineering and Knowledge Engineering*, 17(02):171–190, Apr. 2007.

[6] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[7] E. Kamar, Y. Gal, and B. Grosz. Incorporating helpful behavior into collaborative planning. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 875–882, 2009.

[8] E. Kamar, Y. Gal, and B. Grosz. Modeling information exchange opportunities for effective human-computer teamwork. *Artificial Intelligence*, 2012.

[9] L. L. I. R. on Care Integration. *Order From Chaos: Accelerating Care Integration*. National Patient Safety Foundation, 2012.

[10] A. Rao, M. Georgeff, et al. BDI agents: From theory to practice. In *The international conference on multi-agent systems (ICMAS)*, pages 312–319. San Francisco, 1995.

[11] M. Roth, R. Simmons, and M. Veloso. What to communicate? execution-time decision in multi-agent POMDPs. *Distributed Autonomous Robotic Systems 7*, pages 177–186, 2006.

[12] D. Sarne and B. J. Grosz. Determining the Value of Information for Collaborative Multi-Agent Planning. *Journal of Autonomous Agents and Multi-Agent Systems*, To appear 2012-2013.

[13] T. Wagner, J. Phelps, V. Guralnik, and R. VanRiper. An application view of coordinators: Coordination managers for first responders. In *AAAI*, pages 908–915, 2004.

[14] N. Yorke-Smith, S. Saadati, K. Myers, and D. Morley. The design of a proactive personal agent for task management. *International Journal on Artificial Intelligence Tools*, 21(01), Feb. 2012.