



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Information Sharing for Care Coordination

The Harvard community has made this article openly available.
[Please share](#) how this access benefits you. Your story matters.

Citation	Amir, Ofra, Barbara J. Grosz, Roni Stern, and Lee M. Sanders. 2013. Information sharing for care coordination. In Proceedings of the 2nd International Workshop on Human-Agent Interaction Design and Models: 6-10 May, 2013, Saint Paul, MN, ed. AAMAS, 82-89. IFAAMAS.
Published Version	http://haidm.wordpress.com/2013/05/07/proceedings-online/
Accessed	May 27, 2017 9:08:09 PM EDT
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:10652550
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

(Article begins on next page)

Information Sharing for Care Coordination

Ofra Amir¹, Barbara J. Grosz¹, Roni Stern¹, and Lee M. Sanders²

¹ Harvard School of Engineering and Applied Sciences, Cambridge MA 02138, USA
oamir@seas.harvard.edu, grosz@eecs.harvard.edu, roni.stern@gmail.com

² Center for Policy, Outcomes and Prevention, Stanford University, Stanford, CA
94305, USA
leesanders@stanford.edu

Abstract. Teamwork and care coordination are of increasing importance to health care delivery and patient safety and health. This paper describes our initial work on 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.

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 [12]. This work 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 conditions³. 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.

Figure 1 illustrates the complex environment in which agents supporting care for a child with a complex condition would operate. As can be seen in the figure, the care team 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 significantly 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, as represented by the horizontal lines in the figure showing time periods in which a caregiver is participating in the child’s care. 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 in one place. The figure

³ 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.

also highlights a distinguishing challenge of care for these children: their developmental stages (see horizontal center of figure) affect and may be affected by treatment, adding uncertainty to plan development and typically necessitating plan revision. 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.

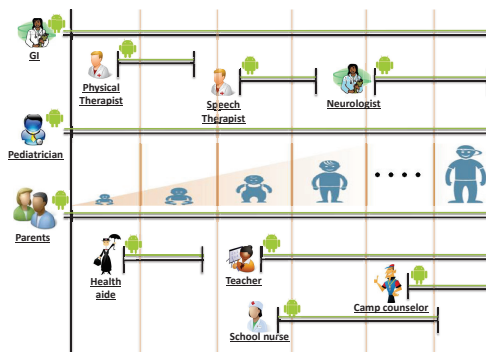


Fig. 1. Agents in the General Care Context

2 Challenges in Health Care Coordination

Our planned 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. They would also track follow-up efforts to help ensure that treatment activities—including those for information-sharing—are carried out. The care for children with complex conditions has several distinguishing characteristics that make plan support challenging, including the following:

An evolving team: The various care providers differ in their expertise, knowledge about a child's condition, and concern with a child's longitudinal care plan. The team changes over time, as new providers may join, existing members may leave, and some may be active only intermittently. These characteristics differ radically from those of prior MAS work that has considered issues of forming teams and developing coordinated or collaborative plans.

Uncertain, evolving action sets: The commonly made “closed world” assumption, i.e., the set of actions and goals are constant over time, does not hold in long-term care planning, as the child changes developmentally over time, and new medical treatments may come into play. Planning needs to accommodate new actions (e.g., new treatments and therapies) and remove actions that are no longer relevant.

Conflicting goals and multiple time scales: As the care plan is executed, conflicting goals may arise, either from limited resources or from contradictory effects of actions at different time scales. For instance, an action might help achieve a short term goal but conflict with a long-term goal. Providers often fail to detect such conflicts until after their impact on a child occurs. Furthermore, in addition to achieving goals specified in the care plan, there are maintenance goals [5] that need to be considered.

The development of agents that are effective in such settings introduces several significant challenges. New **information exchange** capabilities are required for evolving groups to function as a team. Furthermore, as actions are added and removed or when conflicting goals are identified, agents will need to **alert the right set of providers at the appropriate time**. In addition, the changes in the team and action sets often require **re-planning that takes into account the long-term care plan**, which is characteristically different from re-planning for execution failures. 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, as well as for determining care providers’ intended actions and beliefs.

3 Related Work

Prior multi-agent systems efforts e.g., Electric Elves [3], CALO [20], RADAR [7] have addressed the development of multi-agent and planning technologies for personal assistant agents that enable people to better accomplish their tasks in office environments and military settings. Since the primary goal of these projects was to develop a personal assistant agent, they focused largely on the support of a single individual. Some prior work, including Coordinators [19], has addressed collaboration among personal assistant agents, but there are key differences between these efforts and the CASPER goal. First, in prior work, the collaborating agents all operated within a single organization and shared a common vocabulary. The heterogeneity of agents in the health care domain is far greater, as they vary in levels of medical literacy and backgrounds. Furthermore, care givers typically work for multiple organizations. Second, prior work has considered task allocation and the generation of hierarchical plans for agents that, though distributed, are tightly coupled in their efforts [11]. In the health care domain, different providers operate semi-independently and have many competing tasks (e.g. caring for other patients). Their plans are only loosely coupled,

but doing the right thing when they interact is essential to plan success. Furthermore, the underlying assumptions of most of these systems are unlikely to hold for CASPERs given the evolving, longitudinal nature of health care plans as new team members, treatments, actions and even goals are typically introduced to the system and there is the possibility for long-term and maintenance goals to interact complexly with short term goals.

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 [13] 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. Furthermore, they are not suited to incorporate new actions and agents as the care plan evolves, because they assume that a complete model of states and transitions is given in advance and known by all agents. Theories of teamwork and collaboration [8, 4, 17] 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 [10, 15, 14, 6, 18], 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.

4 Approach

Our planned approach will build on the work described in Section 3, as well as that on modeling collaboration [8], helpful behavior [9] and interruption management [10, 16]. 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. [15] to evaluate collaborative agents. In this game, the agents 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.

Recent work has used this domain to evaluate a computer agent interacting with a human teammate [6]. In their setting, the human and agent assume the same role, and are treated as equal collaborators. In contrast, in the healthcare domain, agents are unlikely to replace caregivers in making complex treatment decisions, and therefore we envision the agents as taking on a supporting role and assisting caregivers by reducing the burden of making information sharing decisions. Therefore, we will use a game setting in which human players make

movement decisions and are assisted by agents which decide what observations to share with the other player.

In the second stage, we will evaluate agents in a more complex game setting which has more of the characteristics of the care coordination scenario. To this end we have designed a game in the ColoredTrails framework [1], inspired by discussions with a pediatrician who works with children with complex conditions. Figure 2 shows the game board we have designed. In this game, human players (denoted P1, P2, and P3 in the figure) are located in different positions on the board. There are several goals located on the board, denoted G1 to G4. The players need to agree on a set of shared non-conflicting goals to pursue (goals of the same shape in the figure). This scenario corresponds to caregivers choosing between alternative treatment options in the healthcare domain. There are several sources of uncertainty in the game. First, there is uncertainty about the utility of each goal. In each turn each of the players can move towards one of the goals. Players then obtain more accurate information about the utility of the goal they are approaching. This is analogous to physicians who have some estimate of how well a treatment would work, but can only assess whether it works for a particular patient as they start it. Second, players cannot observe the positions and actions of their collaborating player. This is similar to the healthcare setting in which care providers are typically unaware of all actions performed by others. Players can choose to communicate information to other; however, as in the Colorado/Wyoming game, in this game communicating information is associated with a cost, corresponding to the time required by physicians when they share information or are presented with new information. Finally, we plan to include “distracter” goals for each player. These goals contribute to each player’s personal utility. They are included to make the scenario more realistic by simulating the situation in which each physician not only is concerned with caring for a particular child, but also has additional tasks relating to other patients. The game will be played by people, who would be responsible for choosing movement actions. These players would be supported by agents that make information sharing decisions by reasoning about other players’ beliefs and the importance of sharing each observation. This abstract environment will enable us to test different representations and agent designs for supporting information sharing in various settings.

In both the Colorado/Wyoming domain and the more complex Care Coordination game we will perform extensive experiments to test the developed agents designs, algorithms and representations. These experiments will use different agent designs, as we plan to test several approaches including POMDP based decision making algorithms (e.g., POMDP, Dec-POMDP or NED-POMDP [10]), and representations for shared plans such as probabilistic recipe trees [9]. We will test different settings of support from the agents, including settings in which the agent makes information sharing decisions without consulting the person playing and setting in which the agent only recommends what information to share and the person can reject or accept that suggestion. We will also run experiments varying the number of players to increase complexity. Agents’ performance will

be measured based on the utility achieved by human players assisted by agents as compared to games in which people make both movement and information sharing decisions. In addition, we will measure the time savings for people that are a result of removing their responsibility to communicate with each other.

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. We have been working with a pediatrician collaborator and his team to determine the needs of caregivers and learn what currently available systems do and do not provide. The need for systems to support care coordination was identified through interviews with parents and caregivers. We envision providing each caregiver, including the parents of the patients, an assistive agent. These agents be integrated with the existing electronic medical records systems, to gather the relevant information for the various caregivers and patients' families. In this real-world setting agents' performance will be evaluated according to several measures, such as user satisfaction, time savings and helpfulness of the information provided by the agents.




Fig. 2. A game design to test information sharing agents

5 Conclusion

Teamwork and coordination are crucial for providing high quality care, especially for children with complex conditions which have a large, diverse, team of providers responsible for their care. To work as a true team, caregivers need to be aware of relevant treatments carried by others and of information relevant to their care of the child. In this work we propose the development of agents that support such caregivers by making information sharing decisions, ensuring that caregivers have the information required to generate, monitor and revise a shared care plan. Beyond care for children with complex conditions, such agents have the potential to improve the delivery of health care for many patient populations, especially those in which multiple health issues interact (e.g., for cancer and for the elderly). In addition, similar issues arise for plan support in other

long-term, complex team environments. For instance, relief plans following natural disasters (e.g., the earthquake in Haiti) involve diverse teams including local and international medical staff, social workers, educators and others. Some rescuers are involved for only a short time or intermittently. Furthermore, relief plans operate over multiple time scales (e.g., short term rescue endeavors, longer term re-establishment of educational systems).

Acknowledgments. The research reported in this paper was supported in part by grant IIS-0705406 from the U.S. National Science Foundation. We are grateful to Kobi Gal and Ece Kamar for many helpful discussions.

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. P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial intelligence*, 42(2):213–261, 1990.
5. S. Duff, J. Harland, and J. Thangarajah. On proactivity and maintenance goals. In *AAMAS*, pages 1033–1040, 2006.
6. 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.
7. D. Garlan and B. Schmerl. The RADAR architecture for personal cognitive assistance. *International Journal of Software Engineering and Knowledge Engineering*, 17(02):171–190, Apr. 2007.
8. B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
9. 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.
10. E. Kamar, Y. Gal, and B. Grosz. Modeling information exchange opportunities for effective human-computer teamwork. *Artificial Intelligence*, 2012.
11. F. Oliehoek. Decentralized POMDPs. In *Reinforcement Learning: State of the Art, Adaptation, Learning, and Optimization*. Springer Berlin Heidelberg, 2012.
12. L. L. I. R. on Care Integration. *Order From Chaos: Accelerating Care Integration*. National Patient Safety Foundation, 2012.
13. 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.
14. M. Roth, R. Simmons, and M. Veloso. Reasoning about joint beliefs for execution-time communication decisions. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 786–793, 2005.

15. 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.
16. 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.
17. E. Sonenberg, G. Tidhar, E. Werner, D. Kinny, M. Ljungberg, and A. Rao. Planned team activity. *Artificial Social Systems*, 890, 1992.
18. P. Stone, G. Kaminka, S. Kraus, and J. Rosenschein. Ad hoc autonomous agent teams: collaboration without pre-coordination. In *In Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010.
19. 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.
20. 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.