Collaborative Health Care Plan Support

Ofra Amir, Barbara J. Grosz, Edith Law and Roni Stern
School of Engineering and Applied Sciences, Harvard University, USA
{oamir,elaw,rstern}@seas.harvard.edu, grosz@eecs.harvard.edu

ABSTRACT

This paper envisions a multi-agent system that assists patients and their health care providers. This system would support a diverse, evolving team in formulating, monitoring and revising a shared "care plan" that operates on multiple time scales in uncertain environments. It would also enhance communication of health information within this planning framework. The coordination of care for children with complex conditions (CCC), which is a compelling societal need, is presented as a model environment in which to develop and assess such systems. The design of algorithms and techniques needed to realize this vision would yield agents capable of being collaborative partners in health care delivery broadly as well as in other environments with similar properties such as rescue and rebuilding after natural disasters. This paper describes the key characteristics of collaborative health care plan support, defines a set of essential capabilities for autonomous "care-augmenting software agents", and discusses three major multi-agents systems research challenges that building such agents raises: evolving long-term plan management, enhancing team interactions, and leveraging human computation for care plan customization.

Categories and Subject Descriptors: I.2.11 [Computing Methodologies]: Artificial Intelligence – *Distributed Artificial Intelligence*

Keywords: Multi-Agent Planning; Human-Agent Interaction; Crowdsourcing; Healthcare

1. INTRODUCTION

In this paper, we present a vision of a multi-agent system that assists patients and their health care providers by supporting care coordination and by aiding in the communication of medical information to patients and their families. The envisioned computer agents, which we dub "Care Augmenting Software PartnERs" (CASPERs), would be intelligent, autonomous agents working as a team to support an evolving group of providers with diverse capabilities and perspectives in formulating a shared "care plan" that operates on multiple time scales and in uncertain environments, deploying that plan, and monitoring and revising it as needed.

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.

These collaboration-capable computer agents have the potential to assist healthcare providers in being a true team. Enabled with capabilities to detect points at which health-related communications are problematic, CASPERs could also alert providers or patients to potential misunderstandings, thus improving the likelihood that care plans will mesh and patients will understand them and carry out actions assigned to them. CASPERs raise a variety of long-term challenges for the field of multi-agent system (MAS). Their development requires new methods and algorithms for multi-agent plan management and decision-making, novel deployment of techniques in natural-language processing, and innovative designs for health information systems interfaces. It will also benefit from human computation in tackling problems that machines alone are not yet capable of solving.

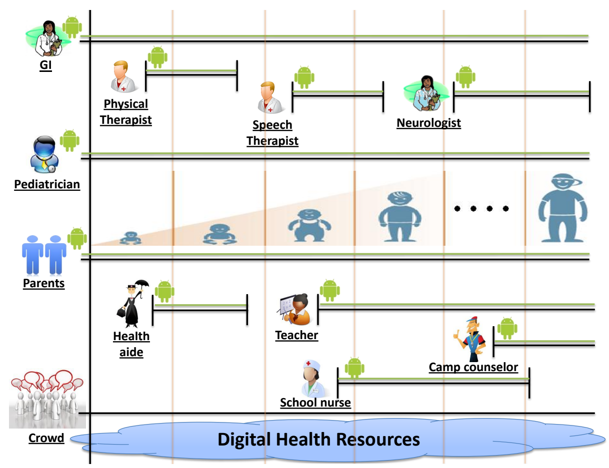

Figure 1: Agents in the General Care Context

We use the care of children with complex conditions (CCC) to illuminate the need for such systems. Figure 1 illustrates the complex environment in which CASPERs supporting care for a CCC would operate. It indicates the potential supporting role of CASPERs as active partners in the team by the green android icons. CASPERs might be instantiated in a variety of ways: as apps on a mobile phone, as agents within a web site, or as enhanced, active EMRs. The figure shows that the care team is diverse and broad in scope, including not only physicians but also other types of care providers (e.g., therapists, teachers). Furthermore, the group of providers may change significantly over time, whether as a result of personnel changes or because the child's condition or developmental stage trigger different needs. Their involvement with the child may be continuous or intermittent, long or short term. Unlike inpatient settings in which the care team meets regularly to discuss a case (e.g., for a cancer patient), the care team for these children seldom all come together in one place. The horizontal center of the figure highlights a distinguishing feature of care for CCCs: that children's developmental stages affect and are affected by treatment, which makes plan coordination and management more complex.

To provide context for describing the challenges of constructing CASPERs, we consider a hypothetical patient Emmy who was born with a progressive neurodegenerative disorder. Among the characteristics of her disorder are an unsteady gait, seizures, and gastrointestinal problems. In addition to her pediatrician and parents, Emmy's care involves a neurologist, GI doctor, nurse manager and their medical staff as well as physical, occupational and speech therapists.

The challenges of integrating long-term developmental goals with consideration of immediate and near-term actions to improve a current condition are evident in the following illustrative scenario: Emmy's parents hope that she soon will be able to use the toilet independently, enabling her to be enrolled in a day care program. At a regularly scheduled appointment, the neurologist recommends a new drug to better control seizures. Unfortunately, the drug causes Emmy to have intestinal problems that make it difficult for her to manage toileting on her own. Had the neurologist known of Emmy's parents' near-term day care goal, he might have suggested an alternative approach.

The medical and health policy communities have recognized the importance of teamwork to the quality of health care and the need to design and monitor "care plans" for CCC [13]. EMRs do not support these activities. There are increasingly many "apps" for chronic care, but they are not well integrated with each other or EMRs. CASPERs also have the potential to improve the delivery of health care for many other patient populations, especially those in which multiple health issues interact (e.g., for the elderly).

There is a compelling need for new kinds of systems, and the MAS community is well-poised to provide important foundations for, and components of, such systems. Addressing the challenges for MAS that CASPERs present also will have impact beyond the health care domain. 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. Furthermore, they may operate over multiple time scales (e.g., short term rescue endeavors, longer term re-establishment of educational systems). Supporting the management of relief plans that are robust and coordinate a diverse team would require many of the same extensions to MAS that CASPERs need.

Existing multi-agent systems, e.g., Electrical Elves [17], CALO [22] and RADAR [5], 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. As the primary goal of these projects was to develop a personal assistant agent, they focused largely on the support of a single individual. While some prior work such as the Coordinators system[18] addressed collaboration among personal assistant agents, there are key differences between these efforts and the CASPER goal. In this prior work, the people whom agents supported all belonged to the same organization and shared a common vocabulary. There is great heterogeneity in the health care domain, and multiple organizations are involved. Another key difference is in the evolving and longitudinal nature of the health care plan, whereas in prior work, plans were intended to be executed within a relatively limited period of time. Lastly, to our knowledge, previous work has not leveraged the wisdom of the crowd to help enhance and customize plans.

The goal of constructing CASPERs to support the care of CCC has obvious, compelling importance in its own right. It can also serve as a generator of exciting and visionary MAS challenges. Subsequent sections describe briefly three such MAS challenges: multi-agent evolving plan management, agent-augmented human interaction and methods for leveraging human computation to empower these agents.

2. EVOLVING PLAN MANAGEMENT

Multi-agent plan support in long-term, complex team environments requires a range of capabilities beyond the current state-of-the-art planning techniques and agent theories. The care for CCC has several distinguishing characteristics 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. For example, as Emmy's disease progresses and her physical condition deteriorates, she may need radically different physical therapy exercises. Even if Emmy's condition remains stable, medical findings may suggest new treatment possibilities.

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, as in the Emmy example in Section 1. Furthermore, in addition to achieving goals specified in the care plan (such as improving mobility), there are maintenance goals [4] that need to be considered, such as maintaining Emmy's ability to properly digest food.

The development of CASPERs that are effective in such settings introduces several significant challenges. New **information exchange** capabilities are required for evolving groups to function as a team. For instance, when Emmy starts attending school, her teachers and the school nurse will join the care team, while her kindergarten teacher leaves it. The teachers need to know the conditions to watch for and the school nurse needs to coordinate any care she might deliver with other providers. Furthermore, as actions are added and removed or when conflicting goals are identified, CASPERs 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.

Existing techniques for collaborative multi-agent planning are not fully able to address these challenges adequately. Classical BDI planning and agent frameworks [15] assume a closed world in which 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 and known by all agents. Theories of teamwork and collaboration [6, inter alia] support collaborative multi-agent planning, but assume a fixed action library and that the group of agents remains fixed after team formation. In contrast, care teams for CCC involve a potentially changing group of providers who interact during different time periods, and it requires planning for both long and short term goals.

3. AUGMENTED INTERACTION

Communication with patients and among providers is essential to the successful execution of health care plans and especially crucial for complex evolving multi-agent plans like those described in Section 2. The effectiveness of such communications varies greatly for several fundamental reasons.

Health literacy: Providers may have inaccurate models of others' expertise or the completeness of their knowledge of a particular case. They may not be calibrated on patients' level of health literacy, which is often low, limiting their ability to follow medical information.

Plan awareness: The usual pressures of immediate care delivery for providers and the sensitive emotional state of patients, may make it difficult for both providers and patients to consider the context of the long-term plan. Such considerations may be crucial to the success of the care plan.

The development of CASPERs offers the opportunity to take prior work on personal assistants [10, 22] to a new level by intervening to augment care-team members' communication with one another or with patients. CASPERs can be helpful bystanders of patient-provider interaction, facilitating patient-provider communication by "whispering in the ear" reminders and suggestions of topics to discuss, a capability we will call **augmenting interactions**. CASPERs might draw on crowdsourcing techniques to retrieve information from the Web or elicit people's help (see Section 4). For instance, Emmy's parents' CASPER might accompany them to appointments playing a role similar to that of a trusted friend or patient advocate. At the neurologist's office, when a new drug is proposed, the CASPER could tap into online databases to check potential conflicts with Emmy's developmental goals, prompting Emmy's parents to ask about side effects should it find any or alerting the pediatrician (or his/her CASPER) and suggesting this new treatment be assessed against the overall care plan.

To perform successfully as helpful bystanders, agents will need reasoning and decision-making capabilities adequate for **aggregating information** from various sources (e.g., medical records, the Web) and determining which information would be helpful [9] with respect to care plan goals, **designing a presentation** appropriate to the context and intended audience (e.g., for email vs. face-to-face, and patients vs. social workers); and determining **when to introduce** the information, which will require not only minimizing the cost of an interruption, but also taking into account costs to interpersonal aspects of the interaction.

A variety of prior work on human-computer dialogue and intelligent user interfaces is relevant to augmented interaction. Dialogue systems research provides models of dialogue structure, coherence and intention recognition [7; 12; 1, inter alia] as well as techniques for modeling the beliefs and intentions of other agents [20; 6; 16, inter alia]. Work in natural-language generation has developed techniques for producing text that is natural and appropriate in context [14; 8, inter alia]. Research in multi-modal interaction and intelligent user interfaces has developed methods for choosing the appropriate medium in which to convey information [19, inter alia]. Recent works have shown the potential usefulness of displaying information simultaneously to physicians and patients [21], and of sharing visit notes with patients [3].

CASPERs will, however, need capabilities significantly beyond current interruption management and belief tracking algorithms and technologies, as well as beyond the state-of-the-art in natural-language generation (NLG), adjustable autonomy and interruption management. Prior efforts have focused on situations in which an agent participates in a one-on-one conversation or performs offline analysis of a conversation. CASPERs, as secondary participants to a conversation, will need to reason about the dialogue and its participants from "outside". They will thus need to track and model the beliefs and intentions of multiple others (the dialogue participants), their past encounters and their roles in the care plan. They will also need to model the effects of interruptions on interpersonal dynamics. For instance, interrupting patients too soon may decrease their confidence or independent thinking about health care issues. Agents may also need to reason about other CASPERs, as the best way to augment a discussion may be to interact with another CASPER rather than directly with a person.

4. HUMAN COMPUTATION

Health coordination technologies need to convey medical information that is accurate, unambiguous and comprehensible as well as to ensure that actions in a care plan can be performed by the designated participant. With the proliferation of crowdsourcing platforms, it is possible to elicit the help of the crowd to customize care plans continuously, filling in the gaps where machines fall short.

There are several ways to leverage human computation [11] to augment care management plans. For example, actions in the care plan are often associated with a set of medical information, which may not be readily comprehensible to users, depending on their educational and cultural backgrounds. A CASPER could leverage a crowd to **translate medical information** into language that can be interpreted correctly by a particular user, by summarizing, paraphrasing and elaborating (e.g., adding definitions for technical terms). The wisdom of the crowd could also make an action more feasible, by **customizing care plans** based on the constraints and preferences specified by the parents or child. For example, a care plan may need to be substantially adjusted based on the form of transportation that a family has available (e.g., car versus public transportation), and the availability of resources (e.g., support groups) near them. There has been recent work on crowd-driven collaborative planning systems [23], which take as input the qualitative and quantitative constraints specified by a user, and leverage a crowd of workers to generate a plan that satisfies those constraints. A CASPER could adapt a similar technique for generating and iteratively refining a care plan, to take into account difficult to quantify constraints and preferences of the patients and parents.

To create such a crowd-in-the-loop care plan customization system, there is a need to address two fundamental challenges in human computation. First, unlike many existing crowdsourcing solutions which assume that each worker is equally competent at all tasks, the challenge here is that a substantial amount of medical expertise may be required (e.g., to transform technical medical documents, or understand the implications of adding or removing action items to/from a plan). We envision that our system would need to draw from a **hierarchical crowd**, comprised of individuals with varying levels of medical expertise (e.g., doctors, residents, medical students, pre-med students, or citizens with no formal medical training). Second, instead of a one-shot process prevalent in existing human computation systems, the process of care plan elaboration needs to involve **iterative feedback**: the requesters (i.e., parents and doctors of the patient) should be allowed ways to monitor the plan as the crowd refines it, and provide feedback to steer the crowd towards certain types of solutions. Together the hierarchical crowd and iterative feedback mechnisms would enable new kinds of joint efforts of people and agent systems, in particular the joint customization of a care management plan. For example, a web-mining system might extract information from the Web or recent medical journals, that is relevant to the patient's questions, and a hierarchical crowd might then evaluate, summarize, translate the information, and incorporate it into the care plan.

5. CONCLUSION

This paper envisions multi-agent systems as collaborative partners to health care providers, supporting both improved care coordination and enhanced communication of health information. The care of children with complex conditions, important in its own right, exemplifies the need for such agents and provides an exciting and challenging environment in which to develop more powerful agent technologies. The paper discusses fundamental advances in a variety of MAS areas which the challenge of realizing this vision would stimulate and which would enhance agent capabilities in ways important for a wide spectrum of problem domains.

6. ACKNOWLEDGEMENTS

This work was supported in part by grant number IIS-0705406 from the U.S. NSF. We are grateful to Ece Kamar and Lee Sanders for many helpful discussions.

7. REFERENCES

[1] R. Barzilay and M. Lapata. Modeling local coherence: An entity-based approach. *Computational Linguistics*, 34(1):1–34, 2008.

[2] D. Bernstein, S. Zilberstein, and N. Immerman. The complexity of decentralized control of markov decision processes. In *UAI*, pages 32–37, 2000.

[3] T. Delbanco, J. Walker, J. Darer, J. Elmore, H. Feldman, S. Leveille, J. Ralston, S. Ross, E. Vodicka, V. Weber, et al. Open notes: doctors and patients signing on. *Annals of internal medicine*, 153(2):121, 2010.

[4] S. Duff, J. Harland, and J. Thangarajah. On proactivity and maintenance goals. In *AAMAS*, pages 1033–1040, 2006.

[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] B. Grosz and C. Sidner. Attention, intentions, and the structure of discourse. *Computatioanl Linguistics*, 12(3):175–204, 1986.

[8] J. Hunter, Y. Freer, A. Gatt, E. Reiter, S. Sripada, C. Sykes, and D. Westwater. BT-Nurse: computer generation of natural language shift summaries from complex heterogeneous medical data. *JAMIA*, 18(5):621–624, 2011.

[9] E. Kamar, Y. Gal, and B. J. Grosz. Incorporating helpful behavior into collaborative planning. In *AAMAS*, pages 875–882, 2009.

[10] E. Kamar and E. Horvitz. Jogger: models for context - sensitive reminding. In *AAMAS*, pages 1089–1090, 2011.

[11] E. Law and L. von Ahn. *Human Computation: Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan and Claypool, 2011.

[12] K. Lochbaum. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572, 1998.

[13] S. M., C. G., F. B., J. C., and R. R. Best care at lower cost: The path to continuously learning health care in america. Technical report, Institute of Medicine of the National Academies, 2012.

[14] O. Rambow, S. Bangalore, and M. Walker. Natural language generation in dialog systems. In *Proceedings of the International Conference on Human Language Technology Research*, HLT, pages 1–4, 2001.

[15] 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.

[16] E. Selfridge, I. Arizmendi, P. Heeman, and J. Williams. Integrating incremental speech recognition and POMDP-based dialogue systems. In *SIGDIAL Conference*, pages 275–279, 2012.

[17] M. Tambe. Electric elves: What went wrong and why. *AI Magazine*, 29(2):23–27, 2008.

[18] 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.

[19] W. Wahlster. *SmartKom: Foundations of Multimodal Dialogue Systems*. Cognitive Technologies. Springer, 2006.

[20] F. Wang and K. Swegles. Modeling user behavior online for disambiguating user input in a spoken dialogue system. *Speech Communication*, 55(1):84 – 98, 2013.

[21] L. G. Wilcox, J. Gatewood, D. Morris, D. S. Tan, E. Horvitz, and S. Feiner. Physician attitudes about patient-facing information displays at an urban emergency department. In *AMIA Annu. Symp. Proc.*, pages 887–91, 2010.

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

[23] H. Zhang, E. Law, R. Miller, K. Gajos, D. Parkes, and E. Horvitz. Human computation with global constraints: A case study. In *CHI*, 2012.