



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

AI Support of Teamwork for Coordinated Care of Children with Complex Conditions

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, Krzysztof Z. Gajos, Sonja M. Swenson, and Lee M. Sanders. 2014. "AI Support of Teamwork for Coordinated Care of Children with Complex Conditions." Paper presented at the 2014 AAAI Fall Symposium Series, Arlington, Virginia, November 13–15, 2014.
Accessed	May 27, 2017 9:08:05 PM EDT
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:14171326
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)

AI Support of Teamwork for Coordinated Care of Children with Complex Conditions

Ofra Amir and Barbara J. Grosz and Krzysztof Z. Gajos

School of Engineering and Applied Sciences

Harvard University

oamir@seas.harvard.edu, {grosz,kgajos}@eecs.harvard.edu

Sonja M. Swenson

Center for Policy, Outcomes & Preventions

Stanford University

sonjas1@stanford.edu

Lee M. Sanders

Division of General Pediatrics

Stanford University

leesanders@stanford.edu

Abstract

Children with complex health conditions require care from a large, diverse set of caregivers that includes parents and community support organizations as well as multiple types of medical professionals. Coordination of their care is essential for good outcomes, and extensive research has shown that the use of integrated, team-based care plans improves care coordination. Care plans, however, are rarely deployed in practice. This paper describes barriers to effective implementation of care plans in complex care revealed by a study of care providers treating such children. It draws on teamwork theories, identifying ways AI capabilities could enhance care plan use; describes the design of GoalKeeper, a system to support providers use of care plans; and describes initial work toward information sharing algorithms for such systems.

Introduction

The coordination of care for children with multiple, simultaneously occurring chronic conditions presents major health-care challenges. As illustrated in Figure 1, these patients require care from a diverse set of providers. In addition to primary care and specialist doctors, care is often provided by physical and occupational therapists, home caregivers, family members and patients themselves, with community agencies and care coordinators providing care-related support. Despite widespread consensus on a need for greater coordination among providers, health care for most children with complex chronic conditions remains poorly coordinated. As a result, they have high rates of unmet health needs and potentially preventable health care crises, and they account for a disproportionate share of health-system use and costs (Leape 2012).

Health care research has shown that effective systems of care for patients with complex chronic conditions require an *integrated care plan* that addresses patient-centered health goals and provides context for treatment decisions (McAllister 2014). In contrast to clinical guidelines that describe treatment plans of individual providers for specific conditions such as asthma or diabetes, such integrated care plans




Figure 1: The complex care environment.

address multiple conditions and incorporate family-centered functional goals. They thus transcend the practice of individual providers. To better understand the barriers to forming and using such care plans, we conducted a study of care providers for children with complex conditions. The study included interviews and observations of representatives of the care team: parents, primary care physicians (PCP), medical specialists, therapists and administrators.

Our study reveals a variety of features of the current care environment that prevent caregivers from acting as a coordinated team and pose challenges to the use of team-based care plans: Team members differ significantly in their conceptions of care plans and goals as well as in their use. At present, goals and plans are typically set by each care team member individually rather than as a team. Even when team plans are set, they are not monitored or updated, which causes the plans to become quickly obsolete given the uncertainty in the outcomes of actions and the evolving condition of patients. Care teams have deficient or inefficient communication mechanisms, leading to problems of missing information and information overload. These findings show that a set of conditions needed for developing and deploying integrated care plans delineated in a recent report by the Lucile Packard Foundation for Child Health (LPFCH) (McAllister 2014) do not hold in current complex care settings.

Drawing on teamwork theories (Grosz and Kraus 1996; Hutchins 1996), we identify teamwork mechanisms lacking in complex care settings. We describe the design of “GoalKeeper”, a system that aims to support care-plan use and care coordination by providing such teamwork mechanisms. Finally, we discuss our initial steps toward designing information sharing algorithms appropriate for these settings.

Barriers to Care Plan Implementation

This section summarizes study findings related to team members' use of plans and to information sharing.

Perspectives on Goals and Plans. The concepts of goals and plans are relatively new for most providers who vary in both their views on and use of them. In current practice, there are typically multiple individual plans formed by different providers, rather than a single shared team plan, and providers are unaware of each other's plans. Moreover, these plans typically address specific medical needs rather than family-centered goals (e.g., going on a family trip).

Moreover, providers' plans differ from each other in many ways: the types of goals include; whether they define goals at all; the timeline of the plan, and monitoring of the plan. For example, therapists define such concrete goals as being able to move from sit to stand and set a clear timeline for these goals. In contrast, specialists often set more high-level goals such as goals related to the willingness or unwillingness of the family to take extreme treatment measures. Some providers define explicit goals, but do not document the actions they plan to achieve them. Other providers specify their plan of action but do not state their goals explicitly. Parents often feel that there is no clear plan or that they need to track multiple different plans and goals.

Information Seeking by Providers. Providers often belong to different organizations, and thus do not always have access to each others' medical notes. More often than not, they do not have the full picture of a child's condition. Parents also cannot easily access their child's full medical record. When team members do have access to the records, they are overwhelmed by the massive amount of information they encounter. Providers said that finding the *relevant* information in the records is hard and that they often miss important information. At the same time, providers said that these records often lack information important to care, in particular information related to social aspects of the family.

Provider-Provider Communication. Providers said that communication with other providers was slow and frequently deficient. They usually communicate through the electronic medical record (EMR) and email. There is no feedback mechanism that enables them to know whether information they sent was seen. Real-time communication is extremely rare and not scalable due to the size of care teams and the many responsibilities of care providers. Thus, parents of children with complex condition often serve as the main information transmitters between care providers. This reliance on parents leads to problems as it depends on parents' memory and understanding of information they received from other providers. Both parents and providers described this situation as problematic.

Parent-Provider Communication. Parents also serve as important sources of information about their child's status. Providers reported that parents provide them with much of the information on which they rely. Parents synthesize and shape information differently depending on the provider with whom they are talking. For instance, they might share different information with the neurologist and the gastroenterologist, shaping their choices by what they think is most closely related to that provider's aspect of care. Parents are

also asked different questions by different providers. For example, parents said that the doctors are usually interested in high-level information such as whether the child is walking, while therapists are more interested in the day-to-day status.

Mechanisms for Successful Teamwork

This section reflects on the study findings with respect to teamwork in complex care and discusses theories that indicate ways technology could support teamwork in the complex care settings. The task of caring for children with complex conditions presents unusual challenges to providers (Press 2014). For most providers, participating on the care team for such patients is an exceptional situation, beyond the typical care they were trained to provide — most of their patients do not have multiple conditions nor require interaction and coordination with so many other providers. Furthermore, the interacting medical conditions of these patients lead to the formation of teams that are themselves complex.

Our study revealed three characteristics of the nature of care teams that underlie many of the care coordination challenges and communication hurdles they face. First, the teams are diverse, including members from multiple organizations who have varying expertise, training, backgrounds and cultures. Second, team members operate on disparate timescales and focus on different aspects of a patient's care. For example, while a neurologist might see patients quarterly and address only their neurological condition, a physical therapist might see them bi-weekly and focus on their mobility. Providers' varied cultures and context of care lead to different information needs and different perspectives on goals and plans as well as uses of these concepts. Third, each provider on the team is typically involved with many different care teams, each with a different team composition, and each operating over a long time period (years). New team members join over time while others leave, so team composition changes over time. Except for the family and in some cases the PCP, team members are not even aware of all other team members.

We have identified two complementary theories of teamwork and collaboration that offer foundations for designing technology to support these complex teams in their development and use of care plans. These theories contain principles and mechanisms relevant to barriers uncovered in our findings and to meeting the aspirations and principles described in the LPFCH report (McAllister 2014).

The first theory, distributed cognition (Hutchins 1996), is rooted in ethnomethodology and sociology, and identifies principles and frameworks for explaining how collaboration arises in practice. The second theory, SharedPlans (Grosz and Kraus 1996), rooted in multi-agent systems research, provides formal specifications for the design of robust agents able to act collaboratively as members of a team.

Distributed cognition theory (Hutchins 1996) (DCT) argues that cognition is not confined to the individual, but rather requires consideration of social and cultural context and can be distributed among a group of people and through time. Distributed cognition requires different cognitive processes and structures from those of cognitive theories of in-

dividual minds, and it necessitates coordination among individuals, artifacts and the environment. DCT describes several successful team practices relevant for care team implementation. Hutchins (1996) discusses *task decomposition* as a powerful mechanism that allows team members to “attend closely to only a limited set of data.” The resulting *loosely coupled* activities can be executed in parallel by different team members without interfering with each other, reducing the need for communication. Hutchins argues, nonetheless, that team members must be *aware of the “big picture”* because while engaged with low-level tasks, they tend to lose sight of the high-level goals, leading to suboptimal behavior.

SharedPlans (SP) (Grosz and Kraus 1996; Grosz and Hunsberger 2006) provides a specification of team plans to provide a framework for effective implementation of multi-agent systems. SP is rooted in the observation that collaborative plans are not simply a collection of individual plans, but rather a tight interleaving of mutual beliefs and (coordinated) intentions of different team members. It specifies the beliefs and intentions required of team members for successful collaborative activities. In particular, the specification stipulates that (1) each team member must commit to the team’s performance of the group activity; (2) team members must establish agreement on a “recipe” for carrying out the group action; SP allows for recipes to be partial, expanded over time, and revised; (3) the group must agree on an allocation of tasks in the recipe according to participants’ abilities to carry them out (i.e., divide work appropriately); (4) team members must commit to performing tasks allocated to them (i.e., adopt intentions to do those tasks); (5) team members must commit to the success of others in doing their tasks (i.e., adopt intentions that their teammates succeed). SP supports loose-coupling of subtasks by requiring that only the team members selected for doing a subtask determine the recipe for that subtask. It explicitly represents the context of the group activity in all subsidiary intentions and (mutual) beliefs, thus providing the “big picture” to team-members.

The DCT and SP teamwork theories suggest several key supporting roles for technology: (1) making the care plan available and visible to all team members, at the appropriate level of detail; (2) enabling care team members to easily adapt and expand parts of the plan, while ensuring they do not conflict with others’ activities; and (3) providing mechanisms to support efficient information sharing with team members based on their roles in the care plan, conveying relevant information about the status of the plan.

GoalKeeper: Supporting Care Teams

Based on our findings, we have undertaken the design of a system, which we call “GoalKeeper” to support the creation, monitoring and revision of team-based care plans. Our work is complementary to research on care coordination and health informatics, e.g., personal health information management (Pratt et al. 2006) and telehealth systems (Darkins et al. 2008). It is also complementary to work on coordinating clinical guidelines (Sánchez-Garzón et al. 2013).

In accordance with the LPFCH report recommendations for care plans and based on the SharedPlans specifications

for collaborative activities, the main constructs in GoalKeeper are *goals*, *actions* to be taken towards accomplishing the goals, and *status updates* for monitoring progress with respect to goals. We next briefly describe the main capabilities of GoalKeeper and design choices informed by the teamwork theories.

Creation of the care plan. GoalKeeper will provide team members a way to collaboratively define the care plan for a child, giving parents a voice in its design. In particular, it will support the specification of goals and actions to be taken to achieve them, and assignment of team members to goals and actions (i.e., task allocation).

Viewing the care plan. To enable the care team to establish *mutual beliefs*, the care plan will be visible to all team members. Team members will be able to choose the level of detail required for them and focus on different parts of the plan, so that they have the right context to support decision-making in different settings. For example, during an office visit GoalKeeper will support providers by presenting the context of the team care plan, which is complementary to the low-level information in the EMR.

Monitoring the care plan. Team members will be able to report status updates related to the completion of actions and the progress towards goals in GoalKeeper. We expect status updates related to progress towards goals to be reported mainly by parents who observe the child more frequently and extensively. This approach accomplishes several goals: (1) it reduces the burden on providers who already need to input data into the EMR; (2) it supports providers in monitoring progress and outcomes of the care plan by presenting patient status information that is unavailable in current practice, and (3) it increases the engagement of parents in the care and has potential to empower them in their discussions with providers by giving them more data about the child’s condition (Suter, Suter, and Johnston 2011).

Expanding and revising the care plan. GoalKeeper will provide a dynamic framework where goals and actions can be easily modified and added by team members. The changes will be conveyed to other relevant team members based on their involvement in the care plan.

Intelligent information sharing. GoalKeeper will identify and deliver such information as status updates and plan changes to relevant team members based on their role in the care plan. We next describe our work towards developing effective information sharing mechanisms.

AI Based Information Sharing

GoalKeeper’s required information sharing capabilities are beyond the current state-of-the-art of AI (Amir et al. 2013). Belief-Desire-Intention (BDI) approaches to multi-agent planning, e.g., STEAM (Tambe 1997), often base their communication mechanisms on theories of teamwork and collaboration (Grosz and Kraus 1996; Cohen and Levesque 1990; Sonenberg et al. 1992). These approaches, however, typically do not reason about uncertainty and utilities. Both are prevalent in the healthcare domain (e.g., uncertainty about treatment outcomes, costs and benefits of treatments). Decision-theoretic (DT) approaches to multi-agent communication typically reason within a POMDP

framework (Roth, Simmons, and Veloso 2006; Pynadath and Tambe 2002; Wu, Zilberstein, and Chen 2011, inter alia). These approaches assume that all possible observations are known in advance and that the team has a joint policy. In the healthcare domain, however, agents must be able to deal with observations that are not anticipated (or knowable) at planning time. Furthermore, care providers' mutual beliefs are limited to high-level goals; information about ways to accomplish subtasks which are typically assigned to a single agent, is only individually known. This way of decomposing tasks reduces the complexity of team planning but poses a challenge for information sharing.

As a first step toward handling information sharing in the face of such partial knowledge, we have defined the information sharing problem that arises in the healthcare domain as the "Single Agent in a Team Decision" (SATD) problem (Amir, Grosz, and Stern 2014) and specified an BDI-DT approach to address it. Informally, SATD can be described as follows: an individual collaborating in a multi-agent team obtains new information, unanticipated at planning time. This (single) agent has incomplete knowledge of others' plans. It must decide whether to communicate this new information to its teammates, and if so, to whom, and at what time. SATD differs from previously studied multi-agent communications problems in that it does not assume complete knowledge of other agents' plans or policies nor that all observations are knowable in advance. It assumes instead that agents have some knowledge of each other's intentions and plans which can be used to reason about information sharing decisions.

"MDP-PRT", our integrated BDI-DT representation, builds on the strengths of each approach. It integrates the Probabilistic Recipe Trees (PRT) representation of an agent's beliefs about another agent's plans (Kamar, Gal, and Grosz 2009) with a Markov Decision Process (MDP). Results from empirical evaluation show that it outperforms the inform algorithm proposed by Kamar et al. (2009) and that it obtains results close to those obtained by an optimal Dec-POMDP policy even though it does not rely on having access to all possible observations in advance.

Conclusion

While team-based care plans have been shown to improve care coordination for children with complex conditions, they are rarely implemented in practice. Based on a study of care teams, we have designed GoalKeeper, a system for supporting care teams. We described our initial work towards developing AI algorithms to support information sharing in complex care teams. Beyond the potential for improving the quality of care for children with complex conditions, the development of such AI methods would advance the current state-of-the-art in planning and information sharing.

Acknowledgements

This research was partially funded by the Nuance Foundation. We also thank CIMIT for seed funding.

References

- Amir, O.; Grosz, B. J.; Law, E.; and Stern, R. 2013. Collaborative health care plan support. In *Proceedings of the 12th international conference on Autonomous agents and multi-agent systems*, 793–796.
- Amir, O.; Grosz, B. J.; and Stern, R. 2014. To share or not to share? the single agent in a team decision problem. In *Models and Paradigms for Planning under Uncertainty: a Broad Perspective*.
- Cohen, P., and Levesque, H. 1990. Intention is choice with commitment. *Artificial intelligence* 42(2):213–261.
- Darkins, A.; Ryan, P.; Kobb, R.; Foster, L.; Edmonson, E.; Wakefield, B.; and Lancaster, A. E. 2008. Care coordination/home telehealth: the systematic implementation of health informatics, home telehealth, and disease management to support the care of veteran patients with chronic conditions. *Telemedicine and e-Health* 14(10):1118–1126.
- Grosz, B. J., and Hunsberger, L. 2006. The dynamics of intention in collaborative activity. *Cognitive Systems Research* 7(2):259–272.
- Grosz, B., and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence* 86(2):269–357.
- Hutchins, E. 1996. *Cognition in the Wild*. Bradford Books. University Press Group Limited, 8 edition.
- Kamar, E.; Gal, Y.; and Grosz, B. 2009. Incorporating helpful behavior into collaborative planning. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 875–882.
- Leape, L. 2012. *Order From Chaos: Accelerating Care Integration*. National Patient Safety Foundation.
- McAllister, J. W. 2014. Achieving a shared plan of care with children and youth with special health care needs.
- Pratt, W.; Unruh, K.; Civan, A.; and Skeels, M. M. 2006. Personal health information management. *Communications of the ACM* 49(1):51–55.
- Press, M. J. 2014. Instant replay – a quarterback's view of care coordination. *New England Journal of Medicine* 371(6):489–491.
- Pynadath, D. V., and Tambe, M. 2002. The communicative multiagent team decision problem: analyzing teamwork theories and models. *Journal of Artificial Intelligence Research* 16(1):389–423.
- Roth, M.; Simmons, R.; and Veloso, M. 2006. What to communicate? execution-time decision in multi-agent POMDPs. *Distributed Autonomous Robotic Systems* 7 177–186.
- Sánchez-Garzón, I.; Fdez-Olivares, J.; Onaindía, E.; Milla, G.; Jordán, J.; and Castejón, P. 2013. A multi-agent planning approach for the generation of personalized treatment plans of comorbid patients. In *Artificial Intelligence in Medicine*. Springer. 23–27.
- Sonenberg, E.; Tidhar, G.; Werner, E.; Kinny, D.; Ljungberg, M.; and Rao, A. 1992. Planned team activity. *Artificial Social Systems* 890.
- Suter, P.; Suter, W. N.; and Johnston, D. 2011. Theory-based telehealth and patient empowerment. *Population health management* 14(2):87–92.
- Tambe, M. 1997. Agent architectures for flexible practical teamwork. *AAAI* 97(1):997.
- Wu, F.; Zilberstein, S.; and Chen, X. 2011. Online planning for multi-agent systems with bounded communication. *Artificial Intelligence* 175(2):487–511.