Exploring Active and Passive Team-Based Coordination

Priya L. Donti and **James C. Boerkoel Jr.**
Harvey Mudd College, Claremont, CA
{pdonti, boerkoel}@g.hmc.edu

Introduction

As human-robot teamwork becomes increasingly common, a key challenge is to fluidly and intuitively coordinate team members' interactions. In this work, we explore two modalities of human-robot coordination: active, where agents intentionally attempt to understand and influence the plans of human teammates, and passive, where agents simply react to their human teammates' varying behavior. In our Productivity and Wellness Pal (PaWPal) project, we seek to develop an agent that actively elicits a teammate's constraints, preferences, and goals in order to nudge them towards better behavior. Conversely, in our Coordinating Human-Robot Teamwork project, we take a distributed approach to scheduling where agents passively adapt to teammates' plan executions. Our research hypothesis is that human-robot coordination techniques will lead to more natural and effective human-robot teamwork if they recognize and respect the inclinations of all teammates.

The Productivity and Wellness Pal

Robots and virtual agents can help humans navigate the complexities of their daily lives by nudging them towards more optimal behavior. However, in order to generate meaningful suggestions, agents must first understand their teammates' motivations.

A complicating factor in human-robot teamwork is that what humans *want* to do may conflict with what they *should* or *must* do. Furthermore, many AI approaches require users to tediously encode their preferences as a quantitative objective function, which does not lend itself to the way humans naturally express their preferences (Pu and Chen 2008). On the other hand, Conditional Preference Networks (CP-nets) compactly encode users' preferences using qualitative, example-based semantics that intuitively align with users' decision making processes (Boutilier et al. 2004). We hypothesize that building upon previous work that combines CP-nets with other models, such as constraint representations (Boerkoel, Durfee, and Purrington 2010), can provide an effective way to help users evaluate tradeoffs between their competing motivations.

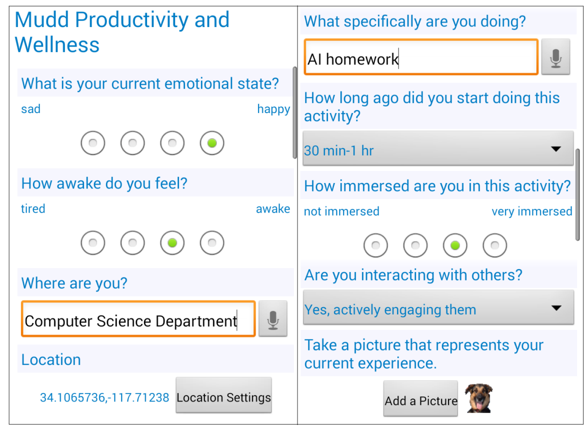

Figure 1: Screenshots from our ESM study, conducted using PACO on Android (http://www.pacoapp.com/).

We introduce our Productivity and Wellness Pal (PaWPal) project as a platform for exploring this hypothesis. PaWPal is a recommendation agent for undergraduate students that seeks to (a) increase human teammates' understanding of what makes them successful, and (b) motivate them towards more optimal behavior by helping them understand the down-stream implications of their decisions. PaWPal is currently in the beginning stages of development as a "virtual" smartphone agent and uses machine learning to characterize its teammate's efficacy at various tasks.

An Experiential Study

In order to understand the needs that PaWPal could fill for an undergraduate teammate, we conducted a study using the Experience Sampling Method (ESM) (Hektner, Schmidt, and Csikszentmihalyi 2007). In ESM, participants are prompted at random times throughout each day to respond to short surveys that capture contextual and situational information about their current experience. ESM is a powerful tool for AI and HRI because it allows us both to understand what needs a particular technology can fill in a user's life and to assess how well a particular technology, once deployed, fills those needs.

Our ESM study lasted for one week in April 2014 at Harvey Mudd College. We signaled twenty-five undergraduate participants eight times randomly throughout their waking hours each day to respond to a short smartphone-based survey about their current activity, emotions, level of immersion, and other experiential aspects. Sample screenshots from this survey are featured in Figure 1. We also asked participants to reflect upon their ESM data via a daily computer-based survey.

From this survey, we gained an empirical understanding of Harvey Mudd students' productivity, wellness, and technological engagement that has informed PaWPal's design. For instance, our preliminary results suggest that the time and location in which activities take place influence students' happiness and effectiveness; students generally reported being happiest and most effective in the afternoon and were more effective in academic campus locations than in residential locations. In addition, participants were generally happier and more effective when engaging actively with others. This empirical data helped us understand that there are indeed concrete features that inform a user's success.

ESM and CP-nets

Participants also reported that they enjoyed responding to ESM surveys, as it helped them become more aware of their emotions and behavior. However, one shortcoming was that the surveys were sometimes overly invasive or occurred when participants could not respond (such as during class). Our hypothesis is that PaWPal can use an adaptive ESM method that adjusts its notification scheme based on a user's calendar and previous ESM responses to predict when the user will be most effective at work, social, or other activities. To test this hypothesis, we will use supervised learning algorithms on data from our ESM study and future adaptive ESM studies to further characterize features important to users' effectiveness. We plan to employ this information to infer a user's implicit "preferences," i.e. the choices that lead to optimal outcomes for the user. We then plan to use these "preferences" to construct a CP-net or similar structure that can be used to generate suggestions and help users understand their natural inclinations. We posit that integrating ESM-like functionality and qualitative reasoning structures into PaWPal will allow it to make insightful recommendations that are grounded in users' actual behaviors as well as heightening users' self-awareness.

Coordinating Human-Robot Teamwork

While PaWPal actively elicits and reasons over a model of a single teammate in order to nudge behavior, our Human-Robot Teamwork project focuses on passively assisting in managing human teammates' activities by building foundational algorithms for a multi-agent setting. These algorithms can be especially helpful in environments where tempo and complexity outstrip people's cognitive capacity to plan optimally (e.g. Pollack 2005; Berry et al. 2011). Past work formally defines the Multi-agent Simple Temporal Problem (MaSTP) for naturally capturing and reasoning over the distributed but interconnected scheduling problems of multiple individuals (Boerkoel and Durfee 2013). As illustrated in Figure 2, MaSTPs can efficiently maintain flexible spaces of scheduling possibilities (Boerkoel et al. 2013),

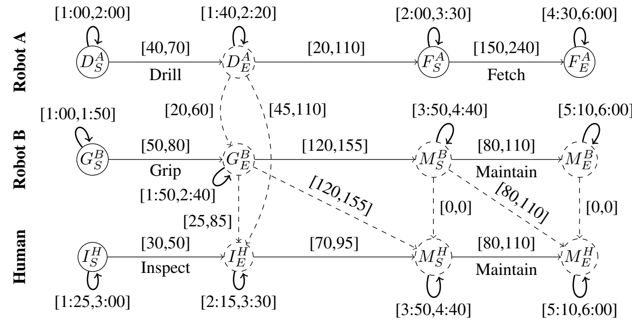

Figure 2: MaSTP representations of human-robot team scenarios provide natural flexibility to compensate for scheduling perturbations (Boerkoel et al. 2013).

which make them well-suited for modeling and adapting to the temporal uncertainty introduced by interactions with human teammates. Past work demonstrates that combining a bottom-up approach—where an agent externalizes constraints that compactly summarize how its local subproblem impacts other agents' subproblems–with a top-down approach—where an agent proactively constructs and internalizes new local constraints that decouple its subproblem from others'—leads to efficient, robust coordination techniques (Boerkoel and Durfee 2013). However, these approaches have only been evaluated in simulation. Our current project looks at empirically evaluating the trade-offs of these various approaches in the context of *real-world* interactive human-robot co-navigation scenarios. Here, agents must coordinate navigating through narrow corridors and timing hand-offs in order to be successful. Key questions include empirically evaluating which approaches (e.g., adaptive scheduling *vs.* pre-negotiated timing of hand-offs) and performance metrics are most indicative of team success.

Discussion

Our work explores a space of approaches that balance between actively shaping human teammates' motivations and gracefully deferring to their tendencies. PaWPal provides a framework for nudging users towards more optimal behavior by eliciting information about their habits and competing motivations. Our Human-Robot Teamwork project, on the other hand, evaluates current adaptive, multi-agent scheduling approaches in robot teamwork scenarios and looks to augment them in response to practical constraints and considerations. Future work includes incorporating insights from both projects to design hybrid approaches that dynamically trade off between active and passive coordination techniques. Further research could also involve designing embodied "study buddies" that help humans make optimal decisions in real-world, collaborative settings. By developing automated scheduling techniques that explicitly account for human inclinations, we hope to improve the coordination between humans and their "teams of technologies."

Acknowledgements

We would like to thank Jacob Rosenbloom and Alex Gruver for their contributions to PaWPal, Bob Evans for his assistance with PACO, and Qutayba Abdullatif and Evelyn Cho for their guidance and assistance with the ESM study. We would also like to thank Zach Dodds, Ben Wiedermann, Emi Reed, and Jeb Brooks for their helpful feedback.

References

Berry, P. M.; Gervasio, M.; Peintner, B.; and Yorke-Smith, N. 2011. Ptime: Personalized assistance for calendaring. *ACM Transactions on Intelligent Systems and Technology* 2(4):40.

Boerkoel, J. C., and Durfee, E. H. 2013. Distributed Reasoning for Multiagent Simple Temporal Problems. *JAIR* 47:95–156.

Boerkoel, J. C.; Planken, L. R.; Wilcox, R. J.; and Shah, J. A. 2013. Distributed algorithms for incrementally maintaining multiagent simple temporal networks. In *Proc. of ICAPS-13*, 11–19.

Boerkoel, J.; Durfee, E.; and Purrington, K. 2010. Generalized solution techniques for reference-based constraint optimization with CP-nets. In *Proc. of AAMAS 2010*, 291–298.

Boutilier, C.; Brafman, R.; Domshlak, C.; Hoos, H.; and Poole, D. 2004. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *JAIR* 21:135–191.

Hektner, J.; Schmidt, J.; and Csikszentmihalyi, M. 2007. *Experience Sampling Method: Measuring the Quality of Everyday Life*. Thousand Oaks: Sage Publications, Inc.

Pollack, M. E. 2005. Intelligent Technology for an Aging Population. *AI Magazine* 26(2):9–24.

Pu, P., and Chen, L. 2008. User-Involved Preference Elicitation for Product Search and Recommender Systems. *AI Magazine* 29(4):93–103.