Advances in Cognitive Systems 4 (2016) Submitted 4/2016; published 6/2016

Improving the Robustness of Team Collaboration through Analysis of Qualitative Interactions

Matthew Klenk KLENK@PARC.COM
Daniel G. Bobrow BOBROW@PARC.COM
Johan de Kleer DEKLEER@PARC.COM
Palo Alto Research Center, Palo Alto, CA, USA

Wendy Mungrovan WENDY.MUNGROVAN@BARNSTORMRESEARCH.COM
Jorge Tierno JORGE.TIERNO@BARNSTORMRESEARCH.COM
Barnstorm Research, Cambridge, MA, USA

Abstract

Members of effective teams must have knowledge about each others future actions. Typically, this is done through messages or precomputed divisions of labor. The former requires ongoing communication between the agents and the latter constrains the autonomy of the individual agents. We introduce *coordination rules* that facilitate collaboration between autonomous agents when communication is lost. By envisioning the results of all possible plan executions for each agent, we identify which decisions result in the greatest increase of within-team uncertainty. If removing this action does not significantly reduce the expected utility of the plan, we create a coordination rule, a statement that the agent will or will not take a particular action in some possible future. Coordination rules facilitate collaboration by improving state estimation and prediction by teammates. To accomplish this, we make the following contributions. First, we identify qualitative interactions by representing the space of decisions made by the agents in *plan-with-options* and their consequences in a *factored envisionment* that compactly represents multi-agent simulations. Second, we define two classes of within-team uncertainty metrics with respect to the envisionment. Third, we present an evaluation of the effects of coordination rules on action selection in three scenarios. In all three scenarios, coordination rules enabled extended planning horizons and reduced planning times with no significant effect on plan quality.

1. Introduction

Coordinated goal-oriented behavior is a hallmark of intelligence. Groups that need to coordinate actions include trained teams (e.g., athletes executing a play), familiar novices (e.g., a family arranging childcare and meals), or complete strangers (e.g., drivers navigating an intersection). Deciding what to do is difficult in these scenarios due to ambiguity introduced by other agents' actions. Coordination can be achieved by constraining autonomy (e.g., athletes executing a play), division of labor (e.g., dad cooks dinner while mom picks up the kids), and social norms (e.g., at a four-way stop, the car on the right gets to go first). In new or changing situations, agents coordinate by considering other agents' goals and reasoning about possible futures. This preserves the ability of the agents to respond autonomously to failures and opportunities that arise during plan execution while still coordinating their behavior, even in absence of communication. We define *coordination rules* as restrictions to the space of actions available to individual agents during execution. We demonstrate that when communication is not possible coordination rules facilitate collaboration by reducing within-team uncertainty.

To enable coordinated behavior, it is necessary to take into account the uncertainty of the dynamics of the world as well as the decisions of other agents. We use a standard planning model to represent the dynamics of the world and introduce *plan-with-options* as a model of contingent plans to represent the space of decisions left to the agent. *Envisioning* from qualitative reasoning [Kuipers, 1994] [Weld & de Kleer, 1989] is a multi-trajectory simulation process that given a model analyzes all qualitatively distinct futures for a scenario. In the multi-agent setting, many of the distinctions captured by traditional approaches are not relevant (e.g., temporal orderings between unrelated events). Therefore, we introduce *interaction-based factoring* that creates a compact representation of possible futures. To automatically construct coordination rules, we define two classes of within-team uncertainty metrics and how they are computed over the envisionment. Finally, we evaluate these ideas over a set of scenarios.

The contributions of this paper include:

- Definition of coordination rules as constraints on an agent's actions to facilitate collaboration
- Exposition of interaction-based factoring algorithm that produces a compact representation of a multi-agent envisionment
- Description and discussion of two classes of within-team uncertainty metrics
- Evaluation of the effects of coordination rules on the action selection process in three scenarios

This is a cognitive systems [Langley, 2012] problem because it focuses on agents representing and reasoning about other agents' behaviors within a planning and execution system. Furthermore, our contributions combine representations and algorithms from artificial intelligence (AI) planning and qualitative reasoning. Finally, we show how an automated analysis of qualitative interactions can facilitate collaborative behavior, an essential trait of cognitive systems.

2. Representing Plans and Futures

Our scenarios involve agents with different capabilities, team-level evaluation, and uncertainty in our own actions, teammates' decisions, and adversary intentions. The Planning Domain Definition Language (PDDL) [McDermott et al., 1998] provides a standard mechanism for describing the dynamics of the world and scenarios. We use PDDL to represent the preconditions and effects of actions. Uncertain effects are modeled using multiple actions with the same name and preconditions but different effects. These actions may be annotated with associated probabilities.[1]

An *envisionment* [Weld & de Kleer, 1989] is a graph of qualitative states with edges indicating the possible successors. When changes are the result of actions, an envisionment is analogous to the state-action space of planning [Ghallab et al., 2004]. States are defined by the collection of propositions that are true and edges in the graph are labeled by the action or set of actions, in the multi-agent case, that change the state.

Consider a scenario with two rovers, a survey rover and an extraction rover, that are collaborating to extract minerals. The world has a base location and four possible mining locations. The survey rover can move between any of the locations and use its sensors to identify if minerals are present. The extraction rover can also move to any location and use its drill. There is also a joint action where the survey rover uses its sensors to guide the extraction rover's drilling action. The full envisionment has 200 states.

In the worst case, the number of states in the envisionment is exponential in the size of the set of ground predicates in the scenario. Therefore, we introduce two concepts that facilitate representing large scenarios.

[1] Probabilistic PDDL [Younes & Littman, 2004] would perhaps be a cleaner representation, but it is less supported by existing planners.

2.1 Plan-with-Options

One cause for the large state-action space is that every afforded action will result in a possible state transition. To reduce this space, we want to consider the agents' decision process, and to accomplish this, we need to have an explicit representation of the options they would consider. Therefore, we define a plan-with-options as the set of all possible sequences of actions an agent could take.

Consider the rover example, in which plan-with-options can be used to constrain the search pattern of the agents. For example, the survey rover may only visit the locations in a particular sequence with the option of returning home after visiting each one, and the extraction rover must follow the same pattern and only drill when co-located with the survey rover. Instead of having to generate successor states from all pairs of rover locations for the extraction rover's drilling action, we only generate successor states for the drilling action from states in which the rovers are co-located.

We represent these choices using hierarchical task networks (HTNs) [Erol et al., 1994]. HTNs represent tasks hierarchically and HTN planning involves decomposing tasks into subtasks and finally primitive actions that can be executed by the agent. An HTN *method* specifies the decomposition of a task. If multiple methods apply to the same task, then there are multiple plans that satisfy the task. We define the *plan-with-options* to be all possible decompositions for a goal task, as well as the current sequence of actions each agent is pursuing. Therefore, instead of reasoning over every possible action available to every agent, we consider the space of decomposition decisions faced by each agent. This provides the agent autonomy to respond to changing conditions while remaining committed to the collaborative activity.

To implement plan-with-options, for each scenario, we define an HTN for each agent. We model options with multiple decompositions for same task. We translate the HTN into PDDL with Alford *et al.*'s approach [Alford et al., 2009] and use SAPA, a domain-independent planner, [Do & Kambhampati, 2003] to generate all possible sequences of states.[2]

2.2 Interaction-based Factoring

The other main driver in the complexity of the envisionment is the explicit representation of temporal differences between agents' actions. If the survey rover is headed toward location 1 and the extraction rover is headed toward location 2, standard representations would include the following states:

1. Both rovers en route to their respective locations
2. Survey rover at location 1 and extraction rover en route to location 2
3. Survey rover sensing at location 1 and extraction rover en route to location 2
4. Survey rover en route to location 1 and extraction rover at location 2
5. Survey rover at location 1 and extraction rover at location 2
6. Survey rover sensing at location 1 and extraction rover at location 2

It does not matter which arrives first, unless they are interacting with each other from their respective locations. Therefore it should be possible to represent the survey rover's actions with three states (en route, at location 1, and sensing at location 1) and the extraction rover's action as two states (en route and at location 2). While this appears to only remove a single state, consider what happens when we allow the extraction rover to drill when it reaches its destination. In the baseline case, there will be three additional states (one for 4, 5, and 6), but, when considered separately, there would only be a single additional state.

[2] This implementation was for engineering convenience. Future approaches should modify an HTN planner directly.

Model decomposition has a long history to improve scaling of qualitative simulations. For example, DecSIM separates the envisionment into components with limited interactions [Clancy & Kuipers, 1997]. This idea was extended to battlespace planning by analyzing the spatio-temporal trajectories of each agent [Hinrichs et al., 2011]. We extend this idea with *interaction-based factoring* in which decisions about when agents are reasoned about individually versus as a group are made using the planning domain model. The core idea is that if an action includes multiple agents in its arguments, then those agents are *interacting* in that action. Factored envisionments consist of *factors*, which are envisionments of a subset of the agents and *factored states* that exclude facts related to other agents. This form of agent-based factoring has been studied in the multi-agent planning community [Brafman & Domshlak, 2008] with the focus on reducing the complexity of planning.

2.2.1 Algorithm

Algorithm 1 outlines our interaction-based factoring approach. It works in two stages, both of which operate recursively. In the first stage, we identify subsets of agents that are able to take actions together in the current state. That is, there exists an action for which the preconditions are met and each of the agents is a participant. For each subset of agents, including the original set, we create a factored state, which includes the facts related to those agents and objects in the world. From this factored state, we create an envisionment using all the actions in which the subset of agents are interacting. This envisionment is a factor. For each state in each factor, if further subsets of agents from that factor are afforded actions, we create additional factors from each of those factored states. In the second stage, we identify sets of factored states from across the factors of the previous stage called *join state sets*. Join state sets include a single factored state from each factor being joined. Each join state set results in a new factored state if there are new actions that can be taken by the set of agents in the join state set. This recursive factoring and joining approach continues until there are no further actions that result in new factors to be created.

Consider our rover example again, the envisionment includes three sets of agents: one for each rover and one for the rovers together. A single trajectory will involve splitting from the initial state as both rovers move to different locations, joining when they collaborate to drill the minerals, and separating again as they travel back to home. The complexity due to the locations of the rovers alone goes from 25 states (one for every combination of rover locations) to 15 states in the factored envisionment (five locations for the survey rover factor, five locations for the extraction rover factor, and five states for each co-location). In the full envisionment, the state changes for sensing and drilling actions are multiplied by 25, but in factored envisionment, the sensing actions are in the survey rover factor and the drilling actions occur in the extraction rover factor. This results in an exponential reduction in qualitative states in the factored envisionment.

2.2.2 Factoring Evaluation

As shown in the example, there is a significant savings when moving to the factored envisionment. Theoretically, if the state-action space of two individual agents are M and N, then their combined envisionment would include $M \times N$ states. If there was no interaction, their factored envisionment would include $M + N$ states. If they were interacting in every state, then the factored envisionment would still be on the order of $M \times N$ states.

Most collaborative problems have a mixture of interaction and independence. To evaluate the impact of factoring on the size of the state graph, we created a series of scenarios with increasing complexity. These scenarios involved a set of aircraft striking a set of defended targets. Scenarios were made more complex by adding additional aircraft, locations, and targets.

Figure 1 plots the problem complexity against the number of states. In the smaller scenarios, all of the agents are constantly interacting and the factored envisionment is slightly larger. This is due to the overhead

Algorithm 1 Factored envisioning algorithm

```
procedure fenvision($S_{init}, D_{full}, Agents$)
    $FS \leftarrow \emptyset$
    for all $D_i, A_i \in D_{full}, A$ do
        $FS_i =$ factor_state($A_i, S_{init}$)
        if envision($FS_i, D_i$) then
            add_child($S_{init}, FS_i$)
            $FS$.add($FS_i$)
            if factorable($A_i$) then
                for all $js \in$ all_successors($FS_i$) do
                    fenvision($js, D_{full}, A_i$)
                end for
            end if
        end if
    end for
    for all $D_i, A_i \in D, A$ do
        for all $jg \in$ possible_join_graphs($FS, A_i$) do
            for all $jss \in$ identify_join_state_sets($JG_i, D_i$) do
                $js \in$ create_joint_state($jss$)
                envision($js, D_i$)
                for all $s \in$ all_successors($js$) do
                    fenvision($js, D_{full}, A_i$)
                end for
            end for
        end for
    end for
```

Figure 1. The effects of factoring on the size of the envisionment on air strike scenarios.

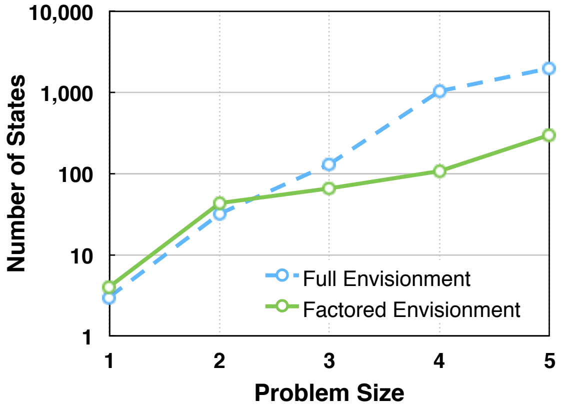

of creating initial states for each factor. As the problems get more complex, there is an apparent order of magnitude savings resulting from reasoning about non-interacting agents independently.

3. Generating Coordination Rules

Coordination rules reduce within-team uncertainty by restricting the space of possible actions for a particular agent. By analyzing the envisionment, the agent measures the amount of uncertainty introduced by each of its decisions. Next, it generates a coordination rule that eliminates the action with the greatest increase in within-team uncertainty in which the remaining decisions still support a successful collaboration. When action probabilities and utility functions are known, we can evaluate the tradeoff between within-team uncertainty and overall expected utility.

To generate coordination rules, it is necessary to have a model of the other agents. We assume that other agents have observations and select actions. While the envisionment captures the space of their possible actions and their outcomes, it does not include any model of their decision-making process. Without this, we assume that each action and outcome is equally likely. Given probabilistic action models and goals, we model the other agents as utility maximizing. Uncertainty in the other agents' observations may be captured in the envisionment (e.g., detection events are modeled as agent actions). Because we cannot ensure future communication, future states of the world are uncertain (e.g., one teammate will not know what the other just did). Therefore, even with a utility maximization model for other agents' actions, there is still uncertainty about what other actions they will take. With or without probabilities, the envisionment captures the space of possible actions for all of the agents. To define within-team uncertainty from an envisionment, it is necessary to answer two questions:

- What aspects of the world are important?
- When do we want to measure uncertainty?

To answer the first question, we identified two types metrics: (1) State-based uncertainty metrics measure the uncertainty of the values of certain predicates in possible future states of the world. (2) Action-based metrics measure the uncertainty of possible action sequences for each agent. For each, it is necessary to determine the time period of interest. For example, when deciding which items to pick up in a grocery store, it is helpful to know what part of the store your partner is in for the next few minutes as it is likely you will run

Scenario	State	Action
SEAD	✓	✗
Patrol	✗	✓
Strike	✓	✓

Table 1. Results for the selection of the expected coordination rule by each metric.

into each other soon. On the other hand, if you are working on a document without communication, you care about which sections of the document your colleague may have changed at the next possible communication point.

Consider a UAV strike scenario in which two UAVs are attacking two targets. We use the following features to create a state-based entropy measure: the location of the UAV, the state of the targets, and the status of the UAV's weapons. We create a unique class for each combination of feature values. To measure the uncertainty at the end of the mission, we assign each end state to its respective class. The distribution of these classes provides a measure of *entropy* [Shannon, 1948]. Coordination rules are considered by ruling out individual actions. For each action, we remove all successor states that require it. From the remaining states, we calculate entropy in the same manner. The difference between this entropy and the original entropy is the uncertainty reduction of removing the action. For the coordination rule, we select the action that results in the largest uncertainty reduction that still allows for successful plan execution. In this case, the UAV 1 tells UAV 2 that if communication is lost, it will not pursue Target 2 while enroute to Target 1.

To better understand the applicability of these types of metrics, we used the following three scenarios with target coordination rules. For the state-based metric, we use the terminal states of the envisionment, and, for the action-based metric, we use the complete action trajectories.

- SEAD : Two aircraft must search two regions and destroy enemy air defenses. In addition to moving actions, there are sensing and attacking actions for both friendly and enemy units.
  - Expected Rule : Aircraft 1 visits target region 1 before target region 2
- Patrol : Two aircraft with different sensors must identify and track targets within an area of interest. Collaboration is required for the search aircraft to hand off targets to the tracking aircraft.
  - Expected Rule : Search aircraft will only follow targets within one of the regions on the return trip.
- Missile Strike : Two salvos of missiles are launched at high-priority and low-priority targets. The targets have countermeasures that may engage the salvos, and the salvos can change what they are targeting.
  - Expected Rule: If each salvo was alive at the point of communication loss, the salvo that was targeting the high-priority target will engage the low-priority target.

Table 1 indicates if each metric was able to identify the expected coordination rule. That is, the metric scores that the action in the coordination rule as contributing the most to within-team uncertainty without impacting the solution quality. State-based uncertainty metrics are most effective when there are clear aspects of the state that are relevant to mission decisions (e.g., the status of a particular target). Action-based metrics work best in scenarios with repeated patterns of activity (e.g., moving between different regions in a patrol).

4. Impact on Action Selection

We claim that coordination rules support autonomous collaboration by facilitating action selection by teammates that cannot directly communicate with one another. To evaluate this, we use decentralized partially observable Markov decision processes (Dec-POMDP) [Bernstein et al., 2000] to formalize the collaborative action selection problem. Coordination rules result in a smaller state-action space, and therefore we expect off-the-shelf solvers to create plans faster and with longer horizons after applying coordination rules. Given that coordination rules remove options, they may reduce the value of the optimal solution. On the other hand, they may also improve the value of approximate solutions by simplifying the action-selection process. We explored these issues over three scenarios.

4.1 Decentralized Partially Observable Markov Decision Process

A Dec-POMDP is defined by a set of agents, a set of states (including an initial state distribution), a set of actions for each agent, a transition function that specifies the probability of each state given a set of actions taken by the agents, a reward function that specifies the reward for being in a state, the set of observations for each agent, and an observation model that specifies the probability of each observation for each agent given a state. A solution is a set of policies, one per agent, which maps each sequence of observations to an action for each agent.

4.2 Method

We used the Multi Agent Decision Problem Toolbox [Spaan & Oliehoek, 2008] (http://fransoliehoek.net/madp) to evaluate the effects of coordination rules. This is an open source toolbox implements a variety of Dec-POMDP solvers and benchmark problems. We selected three scenarios with the following coordination rules added by hand:

- Intersection: Two vehicles at an intersection with actions to move forward or wait. If both cars move forward at the same time, they crash resulting in a negative reward, and if both cars make it through, there is a positive reward.
  - Rule: Car 1 will wait.
- Recycling robots: Two robots either recycle a small amount, a large amount or recharge. While the robots can recycle small pieces independently, they must collaborate to recycle a large piece, which results in the largest reward.
  - Rule: Robot 1 will not attempt to recycle large amounts.
- Missile strike: Same as in Section 3

For each of these scenario, we generated policies for four planning horizons and used two solution methods:

- Optimal - General Multi-Agent A-Star [Oliehoek et al., 2008]
- Approximate - Forward Policy Search with Alternating Maximization [Emery-Montemerlo et al., 2004]

We measured the amount of time required to generate a policy along with the expected utility of the policy.

4.3 Results

Coordination rules provided an order-of-magnitude time savings with minimal effects on solution quality. In each domain, the solver with the coordination rule was able to find a solution for a longer horizon than the uncoordinated case. Furthermore, Figure 2 illustrates that on problems with the same planning horizon that were solved in a non-trivial amount of time, the problems without coordination rules took significantly longer to solve. In the most extreme example, in the recycling robots domain with a planning horizon of four, the solution time with the coordination rule was 22 seconds. Without the rule, it took 86 minutes.

Figure 2. Comparison of planning times with and without coordination rules.

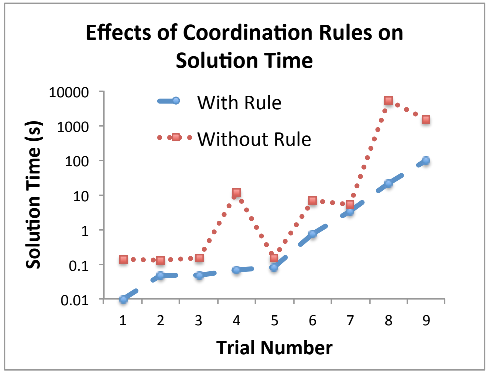

Regarding optimal solution quality, for each planning horizon that generated a policy in each condition, the policy with the coordination rules was on average 3 percent worse than the baseline policy. In the worst case, the coordination rule resulted in a solution that was 6 percent worse for the longest planning horizon on the recycling robots. This is because the coordination rule prevents the highest reward action from being taken.

Regarding approximate solution quality, the results support the claim that coordination rules may result in improved performance when used in conjunction with approximate solvers. In the missile strike scenario, the policies generated by the approximate solution method, were 1 percent better with the coordination rule than those without.

5. Related Work

Most cognitive systems research on collaboration focuses on beliefs and mutual beliefs of the agents. Bello [Bello, 2012] demonstrates the importance of maintaining representations of others beliefs to effectively reason about the world. Kennedy *et al.* [Kennedy et al., 2008] emphasize the importance of learning models of others' capabilities when reasoning about what they might do. Lathrop *et al.* [Lathrop et al., 2011] illustrate how spatial representations provide functional advantages for considering teammates perspectives. Our focus is not on the how others' beliefs are represented, but instead on how an understanding of possible futures can be used to facilitate collaboration. An important extension to our approach would be to include explicit models of others' beliefs into the factored envisionment and explore the implications of such beliefs on the uncertainty metrics.

In addition to cognitive systems, our work builds on research from the knowledge representation, planning and agents communities. An early approach to support multi-agent collaboration involves specifying joint commitments and mutual belief [Cohen & Levesque, 1990]. Joint commitments provide a mechanism for an agent to promise to another agent to either attain a state of mutual belief that the commitment has been satisfied or communicate that it is impossible. The key benefit of this framework is that it prevents agents from dropping out of collaborative behavior as soon as uncertainty arises as they know they will be told by their teammates if they are working toward impossible goals.

Building on this approach is the idea of SharedPlans [Grosz & Kraus, 1996]. SharedPlans were originally developed to further our understanding of dialog systems, but have been extended to be a general model of collaborative behavior [Grosz & Kraus, 1999]. SharedPlans evolve during execution and maintain the mutual beliefs of the collaborating agents and important dependencies between their future actions.

Woolridge and Jennings [Wooldridge & Jennings, 1999] present a definition of the cooperative problem-solving process that goes beyond the focus of this work to include steps for the recognition of a collaborative problem and for team formation. Following the other works, the author's emphasize the importance of joint commitments, but they also introduce the concept of *conventions*, which are understandings within the team when commitments may be abandoned. Our work is complementary to these approaches in that while they facilitate collaboration by agreeing on joint goals and subgoals a priori, coordination rules support collaboration by limiting the action space of the agents.

While the above approaches use logics to support collaborative behavior, others in the multi-agent community view collaboration through mathmatical frameworks describing series of decisions made by individual agents. These include the decentralized partially observable Markov decision processes [Bernstein et al., 2000] and communicative multi-agent team decision processes (COM-MTDP) [Pynadath & Tambe, 2002]. Using a helicopter escort scenario, the authors show how, by analyzing a COM-MTDP, different communication policies can be compared. Our work contributes to these lines of work in two ways: (1) we construct the state-action space through our plan-with-options domain representation and the envisionment process, and (2) we introduce a heuristic that reduces the state-action space by using predictability to limit individual actions. Further study of the cost and benefits of the heuristic in general problems is an important aspect of future work.

6. Conclusion

As cognitive systems become increasing prevalent in our environment, our ability to collaborate with them as well as their ability to collaborate with each other will determine their utility to society. We have shown that plan-with-options provides a natural way to specify the space of coordinated behaviors without unnecessarily removing agent autonomy. In addition, we demonstrated how interaction-based factored envisioning compactly represents the effects of multi-agent plan executions. From this envisionment, we defined two classes of within-team uncertainty metrics, and discussed how they work on different types of scenarios as well as how they enable the selection of coordination rules. Finally, we demonstrated that coordination rules facilitate action selection in three scenarios by removing complexity from the state-action space. This accelerates planning times and allows for longer planning horizons while having minimal effects on plan quality.

The generation and application of coordination rules should facilitate collaboration in many different multi-agent action selection formalisms. Therefore, in addition to exploring different scenarios, future work should address different action selection problem formulations and solution methods. Also, while we have defined two classes of uncertainty metrics, we have not defined algorithms to automate their selection. Furthermore, there are clearly domain specific methods that should also be explored (e.g., remaining near a location that could reestablish communication if lost). Finally, to design systems of collaborating agents, action models will likely exist at multiple scales. In this work, we discuss briefly how qualitative evaluation of state-action space can be used to generate coordination rules that simplify the action selection problem when more detailed models are available. The automatic generation and curation of these multi-level models will be an important problem for long-lived collaborative cognitive systems.

Acknowledgments

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

[Alford et al., 2009] Alford, R., Kuter, U., & Nau, D. S. (2009). Translating HTNs to PDDL: A small amount of domain knowledge can go a long way. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 1629–1634).

[Bello, 2012] Bello, P. (2012). Cognitive foundations for a computational theory of mindreading. *Advances in Cognitive Systems*, *1*.

[Bernstein et al., 2000] Bernstein, D. S., Zilberstein, S., & Immerman, N. (2000). The complexity of decentralized control of markov decision processes. *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence* (pp. 32–37). Morgan Kaufmann Publishers Inc.

[Brafman & Domshlak, 2008] Brafman, R. I., & Domshlak, C. (2008). From one to many: Planning for loosely coupled multi-agent systems. *ICAPS* (pp. 28–35).

[Clancy & Kuipers, 1997] Clancy, D. J., & Kuipers, B. (1997). Model decomposition and simulation: A component based qualitative simulation algorithm. *AAAI/IAAI* (pp. 118–124).

[Cohen & Levesque, 1990] Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial intelligence*, *42*, 213–261.

[Do & Kambhampati, 2003] Do, M. B., & Kambhampati, S. (2003). Sapa: A multi-objective metric temporal planner. *J. Artif. Intell. Res.(JAIR)*, *20*, 155–194.

[Emery-Montemerlo et al., 2004] Emery-Montemerlo, R., Gordon, G., Schneider, J., & Thrun, S. (2004). Approximate solutions for partially observable stochastic games with common payoffs. *Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004. Proceedings of the Third International Joint Conference on* (pp. 136–143). IEEE.

[Erol et al., 1994] Erol, K., Hendler, J., & Nau, D. S. (1994). Htn planning: Complexity and expressivity. *AAAI* (pp. 1123–1128).

[Ghallab et al., 2004] Ghallab, M., Nau, D., & Traverso, P. (2004). *Automated planning: theory & practice*. Elsevier.

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

[Grosz & Kraus, 1999] Grosz, B. J., & Kraus, S. (1999). The evolution of sharedplans. In *Foundations of rational agency*, (pp. 227–262). Springer.

[Hinrichs et al., 2011] Hinrichs, T. R., Forbus, K. D., de Kleer, J., Yoon, S., Jones, E., Hyland, R., & Wilson, J. (2011). Hybrid qualitative simulation of military operations. *IAAI*.

[Kennedy et al., 2008] Kennedy, W. G., Bugajska, M. D., Adams, W., Schultz, A. C., & Trafton, J. G. (2008). Incorporating mental simulation for a more effective robotic teammate. *AAAI* (pp. 1300–1305).

[Kuipers, 1994] Kuipers, B. (1994). *Qualitative reasoning: modeling and simulation with incomplete knowledge*. Cambridge, MA, USA: MIT Press.

[Langley, 2012] Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, *1*, 3–13.

[Lathrop et al., 2011] Lathrop, S. D., Wintermute, S., & Laird, J. E. (2011). Exploring the functional advantages of spatial and visual cognition from an architectural perspective. *Topics in cognitive science*, *3*, 796–818.

[McDermott et al., 1998] McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., & Wilkins, D. (1998). Pddl-the planning domain definition language.

[Oliehoek et al., 2008] Oliehoek, F. A., Spaan, M. T., & Vlassis, N. A. (2008). Optimal and approximate q-value functions for decentralized pomdps. *J. Artif. Intell. Res.(JAIR)*, *32*, 289–353.

[Pynadath & Tambe, 2002] Pynadath, D. V., & Tambe, M. (2002). The communicative multiagent team decision problem: Analyzing teamwork theories and models. *Journal of Artificial Intelligence Research*, (pp. 389–423).

[Shannon, 1948] Shannon, C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*.

[Spaan & Oliehoek, 2008] Spaan, M. T., & Oliehoek, F. A. (2008). The multiagent decision process toolbox: software for decision-theoretic planning in multiagent systems. *Proc. of the AAMAS Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)* (pp. 107–121).

[Weld & de Kleer, 1989] Weld, D., & de Kleer, J. (1989). *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufmann.

[Wooldridge & Jennings, 1999] Wooldridge, M., & Jennings, N. R. (1999). The cooperative problem-solving process. *Journal of Logic and Computation*, *9*, 563–592.

[Younes & Littman, 2004] Younes, H. L., & Littman, M. L. (2004). Ppddl1. 0: An extension to pddl for expressing planning domains with probabilistic effects. *In Proceedings of the 14th International Conference on Automated Planning and Scheduling*.