Advances in Cognitive Systems X (20XX) 1–6 Submitted X/20XX; published X/20XX

Analogical Reasoning for Intent Recognition and Action Prediction in Multi-Agent Systems

Irina Rabkina IRABKINA@U.NORTHWESTERN.EDU
Kenneth D. Forbus FORBUS@NORTHWESTERN.EDU
Department of Computer Science, Northwestern University, Evanston, IL 60208 USA

Abstract

The ability to infer internal states of others is a hallmark of human social reasoning. It is central to both interpersonal interactions and team dynamics. This includes reasoning about relationships between others, such as cooperative and competitive relationships. Here, we present a system that enables virtual agents to exhibit the same social reasoning abilities. In particular, we show that analogical reasoning is sufficient for modeling relationships between virtual agents in stag-hunt, a simple stochastic prisoner's dilemma-style game, without the need for an explicit underlying model of cooperation. Furthermore, we show that the analogical model can predict agents' future actions based on their prior behavior.

1. Introduction

A major component of human social interaction involves making inferences about others' mental states, such as their knowledge and belief states, preferences and desires, and goals and intentions[1]. This is referred to as theory of mind reasoning, and has been widely studied by psychologists (see Wellman, 1992; Carruthers & Smith, 1996). Similar reasoning ability can benefit virtual agents that interact with other agents, whether those other agents be virtual or human. Among other uses, such reasoning can inform an agent's own decision making or allow it to make suggestions to correct another agent's misconceptions or incorrect behaviors (Albrecht & Stone, 2018).

In the context of theory of mind reasoning for virtual agents, the *stag-hunt* game has recently gained popularity. Stag-hunt (Skyrms, 2004) is a multi-player game in which players can choose to pursue a high reward (i.e., a stag) cooperatively or a low reward individually. Recognition of other players' willingness to cooperate is critical to maximizing one's own reward, as successful cooperation (or lack thereof) directly affects each agent's outcome.

The best-known implementation of stag-hunt as a task for situated agents is the Malmo Collaborative AI Challenge, hosted by Microsoft in 2017. The main task of the challenge was a game called Pig Chase, a version of stag-hunt implemented in the game of Minecraft via Microsoft's Project Malmo (Johnson et al., 2016). The goal of agents in the Pig Chase tasks was

[1] Note that while the vocabulary is similar, the rich literature on Belief-Desire-Intention frameworks (e.g., Rao & Georgeff, 1995), which deals with models of the self, is outside the scope of the present work, which deals with models of others.

to maximize their own points in the game. This means that they were situated agents in the Minecraft environment that had to decide whether cooperation would be beneficial based on the other player's actions (see Section 3.2.1 for detailed discussion of cooperation tradeoffs). However, Pig Chase only involved one other agent which operated either randomly or according to heuristic search. Thus, the task could be simplified to a binary classification problem of the other agent, without the necessity of inferring the other agent's intentions or future behaviors.

In fact, the team that won the challenge did so by combining a model of generalized agent type with a Q-learning approach for choosing the best behavior, given the map and the predicted type of the second agent (Xiong et al., 2018). While this approach performed well on the Pig Chase task, it is not clear whether it could extend to more complicated scenarios. For example, it is not clear how the generalized agent type model would hold up against agents with more complex underlying decision algorithms. Furthermore, the Q-learning model needs to be retrained for any new task or domain. Given that even with a task-specific warm start to its policy search, the model required thousands of training episodes for the Pig Chase task, retraining is non-trivial.

Shum et al. (2019) proposed another approach to reasoning about other agents in the stag-hunt game: Bayesian model averaging over an explicit representation of team hierarchies. Unlike the Q-learning model, Shum et al.'s did not require training, but rather the team hierarchy representation and a uniform prior. Because their approach is cognitively inspired, Shum et al. compared their model's performance against human judgements in the same stag-hunt scenarios (see Section 3.1 for the task description) and found strong correlations with human predictions.

We propose another cognitively-inspired model for cooperation recognition and behavior prediction, Analogical Theory of Mind (AToM; Rabkina et al., 2017). AToM has previously been used to model children's theory of mind development and reasoning. Here, we show that AToM can recognize intent to cooperate between players in the stag-hunt game using analogical reasoning. Furthermore, AToM can generate appropriate inferences about future actions of the agents it is reasoning about as a side-effect of reasoning about their cooperation.

We compare our results against Shum et al.'s (2019) on the cooperation recognition task. Because Shum et al. compare their model's performance against human judgments, we do so as well. We find that, while humans are more accurate than both models, our analogical model is more accurate than the Bayesian model at most time steps and overall. Furthermore, our model does not require a pre-defined underlying model of multi-agent cooperation to make its predictions. We also find that our model is able to make reasonable inferences about agents' future actions, but are unable to compare these results to either Shum et al.'s model or human predictions due to a lack of available data.

We begin with a description of AToM and its analogical processes. In Section 3, we describe Shum et al.'s (2019) implementation of stag-hunt and our approach to making the necessary predictions via analogy. We present our results and compare to Shum et al.'s and human performance in Section 4. This is followed by a discussion in Section 6, related work in Section 7, and our conclusions in Section 8.

2. Background

2.1 Analogical Theory of Mind (AToM)

We have previously used analogy to model the development of children's reasoning about others' mental states, called theory of mind reasoning. The Analogical Theory of Mind model (AToM; Rabkina et al., 2017) posits that a combination of analogical processes and feedback (from other humans or experimenters, as in supervised learning) leads to the development of theory of mind reasoning.

We have used AToM to explain how children improve their abilities to reason about others' knowledge and belief states by listening to stories (Rabkina et al., 2017) and by learning certain grammatical structures (Rabkina, McFate & Forbus, 2018). However, the model is not specific to these phenomena. Rather, these are illustrative examples of the kinds of learning that happens every day. As people experience interactions that require reasoning about others' mental states, they learn what inferences and behaviors are appropriate in a given situation. When they act inappropriately due to an incorrect inference, they receive feedback in the form of an upset friend, a confused colleague, or a scolding parent. AToM similarly relies on prediction, feedback, and comparison to improve its theory of mind reasoning.

AToM is implemented as a computation model built on analogical processes. We use the same processes here, as intent recognition and action prediction are a subset of theory of mind reasoning. These analogical processes are described next.

2.2 Analogical Processes

2.2.1 Structure Mapping Engine (SME)

The Structure Mapping Engine (SME; Forbus et al., 2016) takes two structured cases, a *base* and a *target*, as inputs and produces a set of *mappings* between them. Each mapping consists of a set of *correspondences*, a *structural similarity score* which estimates the strength of the mapping, and a set of *candidate inferences*. Correspondences are computed between entities and facts in the base and target. Effectively, two corresponding entities play analogous roles in their respective cases; corresponding facts perform analogous functions. The similarity score takes into account the quantity and quality of correspondences, where quality is measured by the amount of overlapping structure that supports a correspondence. The more a correspondence is supported by higher order structure, the more it contributes to the overall similarity score. Thus, the more deep overlapping structure between cases, the higher the score of their mapping, and the more analogical similarity between them.

Candidate inferences are projections from one case to the other[2] that are supported by the correspondences in the mapping. They represent information that may be true of the case that is being projected onto, but is not explicitly represented. Here, we use candidate inferences from the mapping with the highest similarity score to make both hypotheses about agent cooperation and predictions about their future actions. Cases for comparison are retrieved via MAC/FAC, which we describe next.

[2] Projections from the target to the base are sometimes called *reverse candidate inferences*. Here, we use the general term for inferences in either direction.

2.2.2 Many Are Called/Few Are Chosen (MAC/FAC)

We use the Many Are Called/Few Are Chosen (MAC/FAC; Forbus, Gentner & Law, 1995) model of analogical retrieval to retrieve the most similar previously-observed case. Given a *probe* case and a *case library*, retrieval consists of two stages. During the first stage (MAC), dot products are computed between a content vector representing the probe and content vectors representing each case in the case library. MAC is fast over large case libraries, but its measure of similarity does not account for structural differences. Thus, up to three cases with similar dot product scores are passed from MAC to the second stage, FAC. FAC is slower, but performs structural comparison (via SME) between the probe and each case passed to it by MAC. We use only the top mapping produced by FAC for further reasoning.

3. Experiments

3.1 Task Description

We test the intent recognition and action prediction capabilities of analogical reasoning using the spatial stag-hunt domain from Shum et al. (2019). Stag-hunt was first proposed as an alternative to the prisoner's dilemma set of cooperative/competitive games (Skyrms, 2004). Unlike the traditional prisoner's dilemma, cooperation in stag-hunt does not come with a penalty; if all parties choose to cooperate, each individual wins a greater reward than if they had chosen to compete. However, if an individual chooses to cooperate but their compatriots do not, then the individual receives no reward at all. In effect, stag-hunt is a test of one's ability to recognize others' intent to cooperate.

In a typical spatial stag-hunt scenario, a map with hares and stags is generated. Hares are low-value targets that can be captured by a single hunter without the cooperation of others. Stags, on the other hand, are high value targets that must be captured by a team of (two or more) cooperating hunters. To capture a target, the hunter(s) must occupy the same space as the target at the same time.

In Shum et al.'s (2019) version of the game, three hunters, two hares, and two stags are placed on a 5x7 grid world. Some squares in the grid world are not traversable, creating a variety of spatial layouts across grids. Starting locations of hunters and targets also vary. A stag is considered caught when two or more hunters are in the same square as it at the same time. Similarly, a hare is considered caught when exactly one hunter is in the same square as it.

Shum et al. simulate three time-steps of nine different scenarios. At each time step, each hunter moves 0 or 1 squares up, down, left, or right. Hares cannot move, but stags can move to avoid capture. Thus, no target is ever captured before the third time step, but at least one target is captured in each scenario. In four scenarios only a hare is captured and no cooperation occurs (Figure 1 b, e, f, h). In three scenarios, a pair of agents cooperates to capture a stag (Figure 1 a, c, d). In the remaining two scenarios, all three agents cooperate to capture a stag (Figure 1 g, i).

Shum et al. define two tasks in this domain. The first is team inference: which hunters, if any, are cooperating in each scenario? The second is action prediction: where will each hunter move? All inferences and predictions are made following each time step. Shum et al. propose a model that combines Bayesian inference with an explicit causal model of team hierarchies. They show evidence that their model's results correlate well with human predictions for both tasks. However, data about actual predictions made is limited to line graphs (i.e., no numerical values) for the

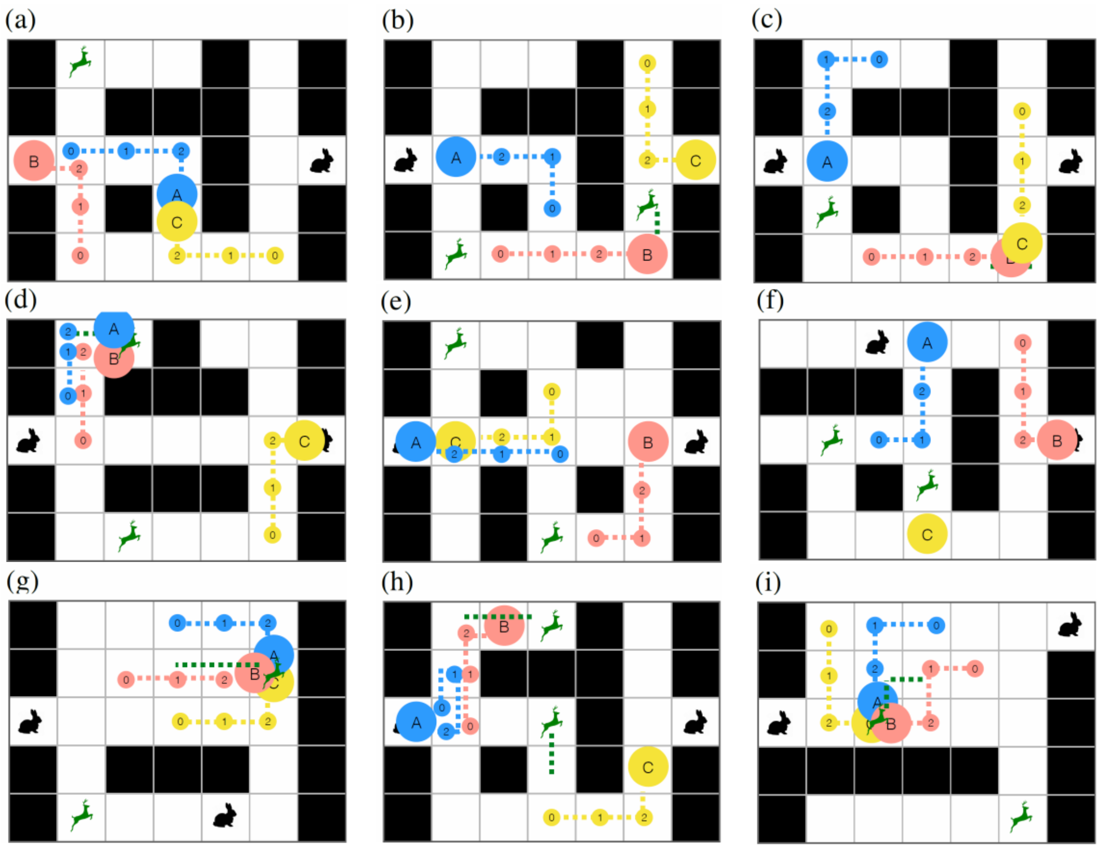

Figure 1. The nine stag-hunt scenarios. Circles represent hunters' position at a given time step. Dotted lines represent motion. Figure adapted from Shum et al. (2019).

team inference task and is not provided at all for the action prediction task. Thus, we compare our model's predictions against ground truth (i.e., actual actions in future steps) for both tasks. For the team inference task, we also compare against Shum et al.'s model and the human predictions, but we note that these comparisons are based on our good faith approximations of Shum et al's results from the figures provided, and may not be entirely accurate.

3.2 Methods

We use analogical reasoning to infer both agents' intent and future behavior. Because SME makes all inferences simultaneously, a single analogical comparison is sufficient to complete both tasks. In this section, we describe case encoding and generation for use with analogy, along with our approach to training and testing for the stag-hunt game.

3.2.1 Encoding and Representations

All scenarios were semi-automatically encoded into predicate calculus from the images provided in Shum et al. (2019). Final spatial representations were based on QSRLib (Gatsoulis et al., 2016). At each time step, we computed (1) whether each individual agent (i.e., hunter or target) is moving, per the *Moving or Stationary* (MOS) library, (2) whether each moving agent has moved closer to or farther from each other agent[3], per *Qualitative Distance Calculus* (QDC; Clementini et al., 1997), (3) whether a pair of agents has overall moved closer to or farther from each other per *Qualitative Trajectory Calculus* (QTC; Delafontaine et al., 2011; van de Weghe et al., 2005), and (4) whether two agents are qualitatively close, far, or located on the same square at all time points (i.e., before step 1, after step 1, after step 2, and after step 3; QDC). Because the agents are situated in a grid world and can move only up, down, left, or right, we used path distance for all distance measurements. Causal relationships between the relations generated in (1), (2), (3), and (4) were also computed (Figure 2).

Non-spatial events (i.e., capture of a target and ground truth cooperation between agents) were manually encoded using the *NextKB* knowledge base (Forbus & Hinrichs, 2017), which integrates materials from several open source ontologies. When appropriate, causal relationships between capture events and cooperative events were also recorded.

3.2.2 Case Generation

Recall that all scenarios in the stag-hunt domain proposed by Shum et al. (2019) include three time steps. Our goal is to make predictions about agents' cooperation and future movements at each step. Thus, we used the representations described above to generate a total of four structured cases for each scenario: three for testing (one per time step) and one for training. The cases used for testing included all computed relations for that time step and all previous time steps. That is, the case for step 1 only had information about step 1, but the case for step 2 had representations for both step 1 and step 2, etc. While capture information was included when capture events occurred, no cooperation events were included in these cases. The training case included information from all three time steps, along with ground truth cooperation events and related causal relations.

a)
```
(causes-PropProp
 (and (holdsIn step1 (approaches agentA agentB))
      (holdsIn step1 (approaches agentB agentA)))
 (holdsIn step1 (closer agentA agentB)))
```
b)
```
(causes-PropProp
 (and (holdsIn step1 (distances agentA stag1))
      (holdsIn step1 (stationary stag1)))
 (holdsIn step1 (farther agentA stag1)))
```

Figure 2. Example qualitative spatial relations between agents in a stag-hunt step. In a), `agentA` and `agentB` (both hunters) move toward each other, resulting in the two agents being closer together than in the previous time step. In b), `agentA` moves away from a stationary `stag1`, causing the two to be farther apart. These causal relationships were computed automatically.

[3] We assume that a given agent does not know where any other agent will move. Thus, closer/farther relations in (2) were computed based on the other agent's location at the previous time step.

a)
```
(and (isa (SkolemFn coop1) CooperationEvent)
     (cooperationParticipants (SkolemFn coop1) agentA)
     (cooperationParticipants (SkolemFn coop1) agentB))
```
b)
```
(and (isa (SkolemFn coop1) CooperationEvent)
     (cooperationParticipants (SkolemFn coop1) agentA)
     (cooperationParticipants (SkolemFn coop1) agentB)
     (cooperationParticipants (SkolemFn coop1) agentC))
```

Figure 3. Two examples candidate inferences for cooperation recognition. a) predicts a cooperation event between `agentA` and `agentB`. It represents zero or one true positive inferences and up to two true negative inferences. b) predicts a cooperation event between all three agents. It represents zero, one, or three correct true positive inferences. Representations are simplified for clarity.

3.2.3 Training and Testing

After all cases were generated, we used leave-one-out cross validation to train and test for each scenario at each time step. This results in a total of 27 testing (1 per time step per scenario) rounds. At each round, we trained using complete cases (including cooperation information), and tested using time step-specific cases (see section 3.2.2, above, for description). As described above, training cases included information about movement at all three time steps, along with ground truth cooperation events. Test cases included only movement information for the time step being tested and any preceding time steps.

During testing, the best overall match for the scenario at the given time step was retrieved. Candidate inferences from the retrieved case to the test case were computed. This resulted in inferences about cooperation events (Figure 3), future movements (Figure 4), and current and future causal relationships. Inferences about cooperation and future movements were automatically identified by their respective predicates.

Consistent with Shum et al. (2019), we analyzed cooperation inferences in terms of hunter dyads. This means that, for each scenario at each time step, up to three cooperation relationships could be inferred.

We similarly used dyads to analyze action predictions. Because we used a qualitative representation to represent movement, all candidate inferences were also qualitative. We compared these inferences to the ground truth at each time step. For example, if the model inferred that `agentA` was moving toward `agentB` at step 3, we checked whether such a motion actually took place at that time step.

a)
```
(holdsIn (SkolemFn step2)
     (approaches agentA stag1))
```

b)
```
(holdsIn (SkolemFn step2)
     (distances agentB agentC))
```

Figure 4. Two example candidate inferences for behavior prediction. These inferences are made during step 1 and makes a prediction about step 2. Representations are simplified for clarity.

4. Results

4.1 Recognizing Intent to Cooperate

4.1.1 Recognition Accuracy

We first compare our model's cooperation inferences to ground truth cooperation events: when two or more hunters captured a stag. Following Shum et al. (2019), we measure accuracy pairwise between hunters, for a total of three predictions for each scenario at each time step (i.e., `agentA` and `agentB`, `agentA` and `agentC`, `agentB` and `agentC`). A true positive inference is one that predicts cooperation between two agents that do, in fact, cooperate. A true negative, on the other hand, is the absence of an inference of cooperation between two agents that do not cooperate in the full scenario. Example candidate inferences are shown in Figure 3.

Table 1 shows our model's overall accuracy. Accuracy is highest (96%) at time step 3, where the model makes only one incorrect inference. At earlier time steps, when less information about agents' behavior is available, accuracy is worse. The lowest accuracy is at time step 1, when 77% of predictions are correct.

4.1.2 Comparison to Human Data and Shum et al. (2019) model

Shum et al. (2019) compare the inferences made by their model against human predictions about cooperation, rather than ground truth cooperative behavior. Both the model's inferences and human judgements are made on a continuous scale that represents degree of certainty that two agents are cooperating. To compare accuracy of our model's predictions to those made by Shum et al.'s model and human participants, we use a 0.5 cutoff. That is, a judgement that cooperation is at least 50% likely corresponds to a positive inference, while a judgement that cooperation is less than 50% likely corresponds to a negative inference. Note, that a lower threshold would correspond to a higher rate of false positives, while a higher threshold would correspond to a higher rate of false negatives.

A comparison of accuracy between human judgement, Shum et al.'s (2019) Bayesian model, and our analogical model is shown in Figure 5. Humans have the highest overall accuracy, tied with the analogical model at step 1 and the Bayesian model at step 2. At step 3, the humans reach 100% accuracy, while the analogical model slightly outperforms the Bayesian.

4.2 Predicting Future Actions

We compare the analogical model's predictions about agents' movement at each future time step to each agent's actual movement at that time step. Because hares are guaranteed not to move, predictions about their movement were always accurate (i.e., the model always predicts that each

Table 1. Accuracy of cooperation inferences made by analogy at each time step.

Time Step	Accuracy
step 1	0.77
step 2	0.81
step 3	0.96
overall	0.85

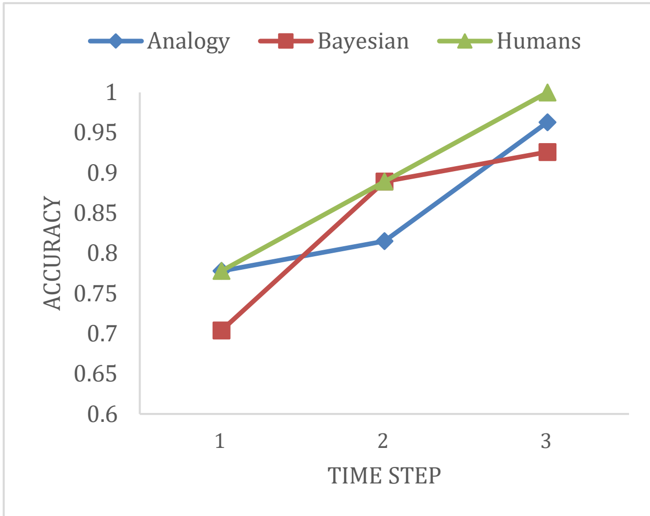

Figure 5. Accuracy of cooperation inferences made by the analogical model as compared to Shum et al's (2019) Bayesian model and human predictions (0.5 probability cutoff). Note that the y axis is shifted to the appropriate range.

hare would remain stationary). We do not include these predictions in our evaluation. Rather, we focus on the movements of the three hunters and the two stags.

For each hunter, a total of six inferences are made per time step (motion in relation to each of two other hunters, two stags, and two hares). For each stag, inferences are only made in relation to the hunters, resulting in 3 inferences at each time step. Three kinds of inferences are possible: `stationary`, `movesToward`, and `movesAway` (Figure 4), resulting in a random baseline accuracy of 33%. The analogical model beats this baseline, with an overall 63% prediction accuracy. Prediction accuracies from time step 1 to steps 2 and 3, and time step 2 to step 3 are shown in Table 2.

5. Discussion

Our analogical model recognizes agents' intent to cooperate in a simple multi-player game better than a Bayesian model and nearly as well as humans. It also beats a random baseline at predicting agents' future behaviors, reaching 63% overall predictive accuracy. These results show that the AToM model, which has previously been shown to accurately model human theory of mind reasoning (Rabkina et al., 2017; 2018) can give virtual agents the ability to perform aspects of theory of mind reasoning—in this case, intent recognition and action prediction—as well.

The model's main task was to infer whether two agents intended to cooperate based on their movements. We assumed that if, at the end of the scenario, two agents cooperated to catch a stag, then they had intended to cooperate at each previous time step. This is consistent with how humans made predictions about the agents' behavior in Shum et al.'s (2019) study: at the final time step, humans inferred that exactly those agents that acted together to catch a stag were cooperating. They gave approximately 100% certainty to cooperation between those agents, and approximately 0% certainty to cooperation between all other agent pairs.

Table 2. Accuracy of action prediction made by analogy at each observation time step for each prediction time step. Note, that no predictions are made at step 3.

Observation Time Step	Prediction Time Step	Accuracy
step 1	step 2	0.61
step 1	step 3	0.64
step 2	step 3	0.64
overall		0.63
baseline		0.33

However, successful cooperation is not the only signal for intent to cooperate, and may, in fact, not be a reliable one at that. We have identified three situations within the stag-hunt game where the assumption that cooperation occurs if and only if it is successful does not hold. The simplest example is *unsuccessful cooperation*. That is, two hunters intend to capture a stag together, but the stag escapes before they are able to corner it. In this case, the intent to cooperate exists, even though the cooperation is not successful.

Similarly, *non-reciprocal cooperation* may occur when one agent intends to cooperate with another, but its would-be partner has other plans. This example points to an inherent flaw in defining intent to cooperate as reciprocal; without access to other agents' internal states (or enough theory of mind reasoning capabilities to infer them), an agent can decide to cooperate with an unknowing partner. It would be incorrect to infer that the two agents intended to cooperate in this case. But it would also be incorrect to infer that no intent to cooperate took place at all. Instead, a unidirectional intent to cooperate should be inferred.

On the other hand, it is possible for two agents to cooperate to capture a stag without intending to. *Non-intentional cooperation* occurs when two agents end up at the right place at the right time. For example, they are both pursuing the same hare and find themselves surrounding a stag. At this point, they might change their plans and decide to capture the stag. Alternatively, it may be that the stag is on both agents' path to the hare. It might be argued that there is an intent to cooperate in first case, albeit only at the last step. In the second case, however, there is no intent to cooperate whatsoever; the hunters capture the stag purely by happenstance. The assumption that cooperation is intended if and only if a cooperative event occurs would lead one to infer that there was an intent to cooperate in both of these cases, including when the agents were, in fact, individually pursuing a hare.

Whether the distinction between two agents cooperating and two agents intending to cooperate matters largely depends on the task at hand. In the present work, where an observer is making inferences about other agents, inferences made with the simplifying assumption may be sufficient. However, if an agent intends to act on its inferences, not considering unsuccessful or non-reciprocal cooperation to be cooperation at all can lead to suboptimal behavior, most likely in the form of missed opportunities to cooperate (and therefore to earn a high reward). Situated agents, then, must have a broader definition of cooperation. Learning such a definition requires feedback based on more than agents' final behaviors. We will consider this in future work.

6. Related Work

Kautz and Allen (1986) described the problem of intent recognition as finding a higher-level explanation for an agent's observed or described behavior. That is, if an agent is acting in order to achieve a goal, another agent should be able to recognize that goal from the set of actions performed. While significant progress has been made in plan, intent, and goal recognition in the intervening years (see Sukthankar, et al., 2014), most methods still require access to the observed agent's plan library, which describes the agent's possible goals and plans (Albrecht & Stone, 2018). When information about internal decision-making is not available, such algorithms cannot be used.

Several approaches have been proposed to avoid this problem. For example, Vered, Kaminka, and Biham (2016) showed that, in continuous domains, online plan generation outperformed plan library-based methods on goal recognition tasks. However, Vered et al.'s approach requires a planner that is capable of performing the given task. Although they use an off-the-shelf planner to demonstrate performance on one task, they build a domain-specific planner for another.

On the other hand, Tambe and Rosenbloom (1994) applied an approach inspired by model tracing to event tracking. In an air combat simulation task, their automated pilot modeled its opponent by following the opponent's steps in a problem-space hierarchy. Tambe and Rosenbloom liken this to pretending to be the opponent, as the goals of the opponent's actions are inferred via very similar mechanisms to those that determine the agent's own actions and goals. The system was implemented within the SOAR architecture and used pre-generated problem-space hierarchies.

Case based approaches to plan recognition require much less task- or domain-specific preparation for successful prediction. Instead, they construct a plan library from observations of agents' actions (Kerkez & Cox, 2003). This allows for more flexible intent recognition, as making predictions about agents in new domains does not require new, domain-specific knowledge. Instead, observed actions and their resultant states are encoded into cases as they are observed. These cases enter the plan library when they differ from previously-stored cases, thus building it on the fly. Although this means that the algorithm cannot perform intent recognition in the first few iterations, its performance improves the more it observes. In Kerkez and Cox's experiments, their case based plan recognizer could make predictions about 60% of planning steps after approximately 8000 observed steps in a sparse domain. In a more concrete domain, the recognizer could make predictions about 90% of planning steps after approximately 20000 observed steps (60% after approximately 2000; 95% after approximately 40000).

Just as having the full plan library of an observed agent is not always feasible, real world applications often require predictions to be made before thousands of planning steps can be observed. Furthermore, real world applications often require reasoning about multiple agents at once, either individually or as a group. While plan recognition approaches have extended into this space (e.g., Zhuo, Yang & Kambhampati, 2012), to the best of our knowledge, they require a full plan library, a complete domain model, or both.

7. Conclusions and Future Work

We have shown that analogical processes can perform both intent recognition and action prediction by observing stag-hunt, a simple multi-agent simulation. However, humans use their social reasoning abilities in much more complex contexts. Furthermore, rather than simply being observers, humans often need to act based on the inferences they have made. To be useful in real-world contexts, virtual agents should be able to do the same. In future work, we will test the social reasoning capabilities of agents equipped with AToM in more complex environments, on more complex tasks, and in situations where the inferences made must drive behavior—and therefore may affect outcomes.

Acknowledgements

We thank Constantine Nakos and Jason Wilson for helpful discussions. This research was supported by the Socio-Cognitive Architectures for Adaptable Autonomous Systems Program of the Office of Naval Research, N00014-13-1-0470.

References

Albrecht, S. V., & Stone, P. (2018). Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258, 66-95.

Carruthers, P., & Smith, P. K. (Eds.). (1996). *Theories of theories of mind.* Cambridge University Press.

Clementini, E.; Felice, P. D.; and Hernandez, D. (1997) Qualitative representation of positional information. *Artificial Intelligence* 95(2):317–356.

Delafontaine, M.; Cohn, A. G.; and Van de Weghe, N. (2011) Implementing a qualitative calculus to analyse moving point objects. *Expert Systems with Applications* 38(5):5187–5196.

Forbus, K. D., Ferguson, R. W., Lovett, A., and Gentner, D. (2016). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, DOI: 10.1111/cogs.12377, pp 1-50.

Forbus, K., Gentner, D. and Law, K. (April-June, 1995). MAC/FAC: A model of Similarity-based Retrieval. *Cognitive Science*, 19(2), 141-205.

Forbus, K. & Hinrichs, T. (2017). Analogy and Qualitative Representations in the Companion Cognitive Architecture. *AI Magazine* 38(4):34-42.

Gatsoulis, Y., Alomari, M., Burbridge, C., Dondrup, C., Duckworth, P., Lightbody, P., ... & Cohn, A. G. (2016). QSRlib: a software library for online acquisition of Qualitative Spatial Relations from Video.

Johnson M., Hofmann K., Hutton T., Bignell D. (2016) The Malmo Platform for Artificial Intelligence Experimentation. *Proc. 25th International Joint Conference on Artificial Intelligence*, Ed. Kambhampati S., p. 4246. AAAI Press, Palo Alto, California USA.

Kautz, H. A., & Allen, J. F. (1986, August). Generalized Plan Recognition. In *AAAI* (Vol. 86, No. 3237, p. 5).

Kerkez, B., & Cox, M. T. (2003). Incremental case-based plan recognition with local predictions. *International Journal on Artificial Intelligence Tools*, 12(04), 413-463.

Rabkina, I., McFate, C., Forbus, K. D., & Hoyos, C. (2017). Towards a Computational Analogical Theory of Mind. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 2949-2954).

Rabkina,I., McFate, C., & Forbus, K. D. (2018). Bootstrapping From Language in the Analogical Theory of Mind Model. *In Proceedings of the 40th Annual Conference of the Cognitive Science Society*.

Rao, A. S., & Georgeff, M. P. (1995, June). BDI agents: from theory to practice. In *ICMAS* (Vol. 95, pp. 312-319).

Shum, M., Kleiman-Weiner, M., Littman, M. L., & Tenenbaum, J. B. (2019). Theory of Minds: Understanding Behavior in Groups Through Inverse Planning. In *Proceedings of AAAI*.

Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.

Sukthankar, G., Geib, C., Bui, H. H., Pynadath, D., & Goldman, R. P. (Eds.). (2014). *Plan, activity, and intent recognition*: Theory and practice. Newnes.

Tambe, M., & Rosenbloom, P. S. (1996). Event tracking in a dynamic multiagent environment. *Computational Intelligence*, 12(3), 499-522.

Van de Weghe, N.; Cohn, A.; De Tre, B.; & De Maeyer, P. (2005). A Qualitative Trajectory Calculus as a basis for representing moving objects in Geographical Information Systems. *Control and Cybernetics* 35(1):97–120.

Vered, M., Kaminka, G. A., & Biham, S. (2016). Online goal recognition through mirroring: Humans and agents. In *The Fourth Annual Conference on Advances in Cognitive Systems*.

Wellman, H. M. (1992). *The child's theory of mind*. The MIT Press.

Xiong, Y., Chen, H., Zhao, M., & An, B. (2018). HogRider: Champion Agent of Microsoft Malmo Collaborative AI Challenge. In *Thirty-Second AAAI Conference on Artificial Intelligence.*

Zhuo, H. H., Yang, Q., & Kambhampati, S. (2012). Action-model based multi-agent plan recognition. In *Advances in Neural Information Processing Systems* (pp. 368-376).