

Najla Ahmad and Arvin Agah*

Plan and Intent Recognition in a Multi-agent System for Collective Box Pushing

Abstract: In a distributed multi-agent system, an idle agent may be available to assist other agents in the system. An agent architecture called intent recognition is proposed in this article to accomplish this with minimal communication. To assist other agents in the system, an agent performing recognition observes the tasks other agents are performing. Unlike the much-studied field of plan recognition, the overall intent of an agent is recognized instead of a specific plan. The observing agent may use capabilities that it has not observed. In this study, the key research question is: What are intent-recognition systems and how can these be used to have agents autonomously assist each other effectively and efficiently? A conceptual framework is proposed to address this question. An implementation of the conceptual framework is tested and evaluated. A set of metrics, including task time and number of communications, is used to compare the performance of plan recognition and intent recognition. This research shows that under certain conditions, an intent-recognition system is more efficient than a plan recognition system.

Keywords: Multi-agent systems, plan recognition, intent recognition, distributed systems.

*Corresponding author: **Arvin Agah**, Electrical Engineering and Computer Science, University of Kansas, 1520 West 15th Street, Lawrence, KS 66045-7621, USA, Tel.: +1-785-864-8821, e-mail: agah@ku.edu

Najla Ahmad: Electrical Engineering and Computer Science, University of Kansas, 1520 West 15th Street, Lawrence, KS 66045-7621, USA

1 Introduction

In a distributed multi-agent environment, agents complete tasks that are assigned in real time or scheduled ahead of time. However, there may be times when an agent may be idle, such as when it has finished its assigned tasks. If instead of remaining idle, this agent assisted another agent in the system during this time, then the task of the other agent would take a shorter amount of time, thus increasing the efficiency of the overall multi-agent system. As an example of a context where this may happen, one can consider the case of robots whose job is to stack boxes. Once a robot has finished stacking its boxes, it will sit idly until assigned another task. A more effective system would have agents that recognize that they have idle time that can be utilized to assign themselves additional tasks. For example, in the above scenario, the robot that has finished stacking boxes can then assist other robots in stacking their boxes. Hence, all of the boxes would be stacked in a shorter amount of time, leading to a more efficient system.

A key component to a multi-agent system is the mechanism that allows agents to interact. In some cases, the agents may be adversaries and pursuing their own individual or team goals. Meanwhile, in a cooperative environment, agents may be trying to achieve an increase in the overall utility of the system. The much-researched idea of plan recognition can be used in either of these situations. In plan recognition, an agent attempts to determine the plan that another agent is following. A plan is a set of actions that lead to a result. Plan recognition is used in situations where communication between agents is to be avoided. An example of this is an operation where communications may be intercepted by an enemy agent.

Expanding on the idea of plan recognition in a cooperative environment, this article proposes a novel concept called *intent recognition*. In plan recognition, an agent observes another agent to determine what plan, or set of actions, that it is currently following. The goal of an agent performing intent recognition is to aid the other agents in the system. By intent recognition, we mean that an agent recognizes what another agent is trying to achieve instead of merely observing the specific steps that the observed agent is following.

It does this by attempting to follow a plan that has the same intent, or overall goal, as the other agents in the system. Similar to plan recognition, intent recognition aims to minimize the amount of communication between agents. Agents that can perform intent recognition are able to autonomously determine that tasks they should be performing in their idle time to increase the overall utility of the system.

An example scenario is that of a person carrying a heavy bag into an apartment building. The plan is to walk to the door of the apartment building, place the bag on the floor, unlock the door with a key, open the door, prop the door open with something, pick up the bag, walk through the door, place the bag down again, and finally shut the door. This plan would take less time if the person's acquaintance was inside the apartment building and saw the person coming through the window. The acquaintance could open the door. The person's new plan would be to walk to the door of the apartment building and walk through the door. Because of the assistance of the acquaintance, the overall task time and the number of plan steps were reduced. Another aspect to consider is that there was no communication in this case. Perhaps the person carrying the bag did not have a free hand to signal to the acquaintance or the acquaintance could not hear through the door. In either case, communication was not possible or necessary for intent recognition to take place.

1.1 Advantages of the Approach

The significance of this research is that it expands the notion of plan recognition to incorporate a new construct, namely intent recognition. It is hypothesized that systems with intent recognition will perform "better" under certain prespecified conditions. As such, more tasks can be accomplished in the same amount of time. As an example, if an agent is pushing a box towards a specified location and it has obstacles in its path, an idle agent that is utilizing intent recognition may then determine that removing all obstacles between the box and the final location will speed up the completion of that task. It would then perform the actions required to assist the other agent, thus decreasing the amount of time required to complete the task.

Another advantage of intent recognition is a reduction in the amount of communication needed for agents to communicate. At any given time, a large percentage of agents in a system can be idle. If all of these agents filled the communication channels with requests for information, communication that is vital to the system may not reach the recipient in a timely manner. Reducing communication results in the reduction of all associated costs.

To study intent recognition, we developed a framework of a multi-agent system in which agents are able to recognize the intent of other agents and utilize their own idle time to assist other agents. This research shows that under certain conditions, an intent-recognition system is more efficient than a plan recognition system.

1.2 Applications

Incorporating intent recognition in future multi-agent systems will lead to agents that are more efficient, have fewer dependencies on human interaction, and are one step closer to accurately emulating their human counterparts. The intent-recognition framework is applicable in many domains. An example of this would be in robotics, particularly if the robots are located in space. An example is a situation where robots are cooperatively working on a task in space, such as space station assembly. There may be cameras in certain locations to monitor the work, but people on Earth may not have a way to view every robot simultaneously. Because of this, people may not be able to determine if robots are working inefficiently at their tasks, are in need of assistance, or are in need of repair. This is particularly true if a robot's communication systems are damaged and it is not able to communicate with people on Earth or with the surrounding robots. If the robots in this cooperative system are able to perform intent recognition, they would be able to determine what the damaged robot is attempting to achieve. These intent-recognition robots would then be able to aid the damaged robot in its task or even repair the robot, depending on their capabilities.

2 Background and Related Work

The study of multi-agent systems is a broad and diverse field [24, 34], with numerous applications, for instance, a sensor web to determine if a weather phenomenon is occurring or not [32]. The key elements of a multi-agent intent-recognition system that we will consider in this article include plan recognition, modeling agents, plan representation, time, communication decisions, and intent.

2.1 Plan Recognition

Plan recognition is used in a system where an agent is observed to determine what series of actions it is performing. Examples include observing to predict agents' destination [16], use of hidden Markov models [23], surveillance purposes [19], soccer simulation [14], modeling opponents in games [9, 28], model building [2, 3], and plan recognition's algorithmic complexity [37].

2.2 Modeling Agents

When playing a game against an opponent, knowing what the opponent intends to do is obviously advantageous. Communication is not an option in this case because intelligent opponents will not willingly reveal their strategies. These include Networks of Influence Diagrams [9], Probabilistic Hostile Agent Task Tracker [10], aiding in the task of learning agents [28], different agents having differing beliefs [25–27], and an agent using its own state as the model [18].

2.3 Plan Representation

A plan is a list of actions followed by an agent to reach a goal, and there are many ways to represent a plan. There are choices on how to represent plans, including the Intelligent Portable Activity Recognition System [19], activities in the RoboCup domain [14], classifying events as complex and simple [13], cost of maintaining and analyzing a list of possible plans [10], reducing the complexity of plan recognition [15], dealing with every previous action of the observed agent having to be witnessed [8], using team plans in plan library [4], and determining when two plans are equivalent [11].

2.4 Time

When events are observed, they can be given temporal relationships. A single- or multi-agent belief network can then be constructed to reflect the temporal nature of the observed actions. Examples include a football game [17], a multi-agent belief network [31], measuring the duration of activities [19], and temporal constraints [14].

2.5 Communication Decisions

The idea of reducing the amount of communication has been explored. These efforts include Behavioral Implicit Communication [6], cooperating agents trying to maximize their overall utility [36], decisions being made without communication [15], communication strategies [22], cost of communication [5, 35, 36], dialogue planning agents [33], dialogue lengths [26, 27], agents deciding when to remind other agents [18], determining when another agent may need assistance [23], and agents recognizing the duration of other agents' activities [3].

2.6 Intent

The subject of intent has been broached in multi-agent systems before, but the research focuses on the domain of story generation. The focus of this type of research is plan generation. The meaning and purpose of “intent” in this case is different from intent recognition in multi-agent systems. An example is where intent is thought of as the intent of the character in the story [30].

Work has been done in the area of human workgroups, which are similar to multi-agent systems, where the coordination takes place between people instead of agents [12, 21]. It has also been shown that research in human work groups is applicable to robotic communities and multi-agent systems [1]. This type of research does not apply to intent recognition because the primary way for humans to convey information to each other is via communication, implicit and/or explicit.

As the provided review of the literature indicates, several frameworks for multi-agent systems have been proposed, such as plan-recognition and inverse models. However, none of these frameworks incorporate the notion of intent recognition that we propose in this article. Intent recognition is the process of recognizing another agent’s objective while minimizing communication. This is different from plan recognition because determining and predicting individual plan steps are no longer the focus of the recognition. This research expands on the idea that agents with differing capabilities can use their individual strengths to work together to solve a common problem. We propose the idea that incorporating intent recognition will enhance the effectiveness and efficiency of a multi-agent system. In this article, we develop a conceptual model based on our intent-recognition framework.

3 Intent Recognition

The greatest advantage of intent recognition over plan recognition is the ability to dynamically choose actions to assist other agents in the system. An intent-recognition agent that is pulling from its plan library may choose an action that it has not observed. Both types of agents, plan recognition and intent recognition, have plan libraries. However, intent-recognition agents have the advantage of not needing the same plan libraries as the agents that are being observed. The general structure for plan and intent recognition is shown in Figure 1.

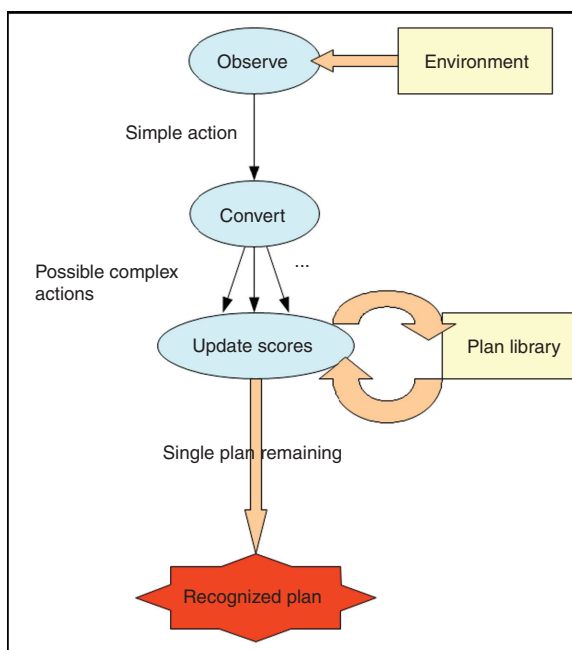


Figure 1. Recognition Structure Overview.




Figure 2. Plan Recognition Process.

A key component to a multi-agent system is the mechanism that allows agents to interact. In some cases, the agents may be adversaries and pursuing their own individual or team goals. Meanwhile, in a cooperative environment, agents may be trying to achieve an increase in the overall utility of the system. The much-researched idea of plan recognition can be used in either of these situations. In plan recognition, an agent attempts to determine the plan that another agent is following. A plan is a set of actions that lead to a result. Plan recognition is used in situations where communication between agents is to be avoided. An example of this is a military operation where communications may be intercepted by an enemy agent. An overview of the plan recognition process is shown in Figure 2.

Expanding on the idea of plan recognition in a cooperative environment, we propose the utilization of the construct of intent recognition in multi-agent systems. The goal of an agent performing intent recognition is to aid the other agents in the system. The agent does this by attempting to follow a plan that has the same intent, or overall goal, as the other agents in the system. Similar to plan recognition, intent recognition aims to minimize the number of communications between agents due to the cost associated with such communication. An overview of the intent-recognition process is shown in Figure 3.

Plan recognition attempts to determine the plan, or set of steps, that an observed agent is following, whereas intent recognition attempts to determine the intent, or overall goal, of the observed agent. The agent performing intent recognition makes observations about another agent in the environment. Instead of trying to find an exact match between the observations and a plan in the plan library, the agent is attempting to find a plan that is similar to the steps that are being observed. Plans are considered similar to the observations on two criteria.

The first consideration is the number of observations that appear in a particular plan. A plan where 75% of the observations occur is less similar than a plan where 90% of the observations occur. When implementing intent recognition, the second plan would have a higher score. The second factor when determining intent is the number of actions in the plan that has not yet been observed. This would also affect the score of the plan when implementing intent recognition.

As an example, there are 10 actions in the system. They are denoted by the numbers 1 through 10. A plan is represented as a list of these actions, for example, {1, 2, 3}. The intent-recognition agent’s plan library would include Plan A={1, 2, 3, 4, 7}, Plan B={6, 7, 8}, and Plan C={1, 4, 7}. The observations made by the agent are the following: Time Step 1, 4; Time Step 2, 7; and Time Step 3, 9. In this case, Plans A and C are more similar to the observation because they both include two of the three observed actions. The base scores of Plans A and C are higher than Plan B. Plan C is chosen as the plan with the most similar intent because there

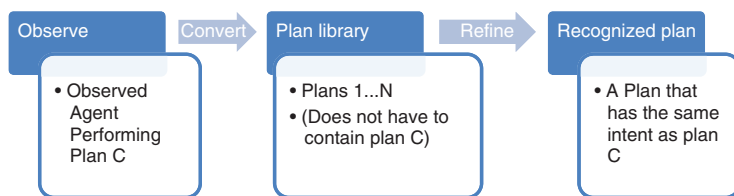


Figure 3. Intent Recognition Process.

is only one action in Plan C that has not yet been observed, whereas there are three actions in Plan A that have not been observed. In this case, Plan C would have the highest score of the plans in the plan library when performing intent recognition.

Plan recognition would look at the above example in a different manner. The third observation of nine would be seen as an outlier to all of the given plans and would reduce their overall scores. After the three observations shown above, plan recognition would not have enough information to recognize a plan. Plans A and C would have the same number of observations corresponding to them. Even if the next observation was 1 and all three steps (1, 4, and 7) from Plan C were seen, these three steps are also seen in Plan A. Additional reasoning capabilities are needed to distinguish these two plans.

4 Research Methodology

We designed and built a collective box pushing simulation environment to test the intent-recognition concept. The environment was developed using the Repast Symphony plugin [29] for Eclipse [7], which is a cross platform Java-based modeling system that facilitates the development of models of interacting agents. The simulated world consists of a 50×50 grid, where boxes are represented by squares, agents are represented by circles, and obstacles are represented by triangles. Agents can move in one of the four cardinal directions. Agents are also able to push boxes in one of the four cardinal directions if they are adjacent to the box and lined up to face the intended direction. If a box or agent encounters the edge of the grid, it stays in place until it moves or is moved in another direction.

Agents are tasked with moving one or more specific box color groups to a given side of the grid. An example of this would be a plan where all yellow boxes are to be moved to the east. There were three types of agent groups (no recognition, plan recognition, and intent recognition). All agents work simultaneously to complete tasks in a cooperative manner. The groups of colored boxes and varying goal direction were used for recognition purposes. In other words, an agent performing recognition had to determine which color of boxes was being moved and in which direction. An illustration of the grid can be seen in Figure 4.




Figure 4. Simulation with Red, Blue, Yellow, and Green Boxes. Agents performing intent recognition are in purple. Observed agents are gray. Triangles represent obstacles.

For each data set, three groups of experiments were done: no recognition, intent recognition (observed and intent-recognition agents), and plan recognition (observed and plan-recognition agents). For each, the experiments were run with agent populations of 5, 10, 15, 20, 25, and 30 agents. The experiments were run five times for each agent population size. There were 11 data sets (obstacle delay, start energy, start plan, boxes per color, and number of obstacles), resulting in a total of 990 experimental runs. The exact values for the lower and upper range of variables were determined via preliminary testing of the system.

The idea of obstacles was introduced into the simulation environment to compare the differences between plan recognition and intent recognition. The number of obstacles in the environment varied from one data set to another. The starting location was randomly assigned by the system at the beginning of each simulation. A snapshot of the simulation is shown in Figure 4.

When an agent encounters an obstacle, it takes a certain number of time-steps to traverse over it. This number was referred to as the obstacle delay. An agent that attempts to cross an obstacle is hindered in two ways. The most obvious way is the time taken to complete a task. However, it is also a drain on energy. The longer an agent stays in one location, the more energy is used up there. In the case where an agent is pushing a box while crossing an obstacle, more energy is used by the agents.

Three types of agents are considered: observed agents, plan-recognition agents, and intent-recognition agents. Observed agents are assigned a plan to follow at the beginning of the experimental run. They have no reasoning capabilities. The observed agents continue following the plan until they either run out of energy or the task is completed. For example, an observed agent can be given a plan to push all of the yellow colored boxes on the grid to the east. The agent will continue this until either it runs out of energy or there are no longer any yellow boxes that are not on the east side of the grid.

Plan-recognition agents begin by choosing an observed agent on the grid to observe. The plan-recognition agent then makes and stores observations of that agent to determine which plan is being executed. Once the plan is determined, the plan-recognition agent then begins executing that plan. It continues until all plan objectives have been met or it runs out of energy. If all plan objectives have been met and the agent has energy remaining, it will begin the observation process again. Intent-recognition agents also begin by choosing an observed agent on the grid. The intent-recognition agent then makes observations until a plan with similar intent is determined. This plan with similar intent is then executed until all plan objectives have been met or the agent runs out of energy. If all plan objectives have been met and the agent has energy remaining, it will begin the observation process again.

Plan recognition differs from intent recognition in that a plan-recognition agent attempts to determine the exact plan that the observed agent is executing. Intent-recognition agents attempt to find a plan in their plan library that has the same overall intent as the plan that the observed agent is executing. The plan determined by the intent-recognition agent may or may not contain actions that the intent-recognition agent has not observed the observed agent executing. Each plan-recognition agent in the system has access to a copy of the same plan library. Each plan is assigned a score of zero at the beginning of the simulation. When an observation is made by a plan-recognition agent that coincides with a given plan, the plan score is increased for that agent's copy of the plan. If the observation does not coincide with the plan, the plan score is decreased.

A single observation can increase the score of multiple plans. For instance, an observation is made that the observed agent moved north. In this example, the plan-recognition agent's plan library consists of three plans:

- GREEN EAST={MOVE TO A GREEN BOX, PUSH THE GREEN BOX EAST}
- GREEN WEST={MOVE TO A GREEN BOX, PUSH THE GREEN BOX WEST}
- BLUE EAST={MOVE TO A BLUE BOX, PUSH THE BLUE BOX EAST}

The plan-recognition agent then observes the environment. If there is a green box to the north, the scores for GREEN EAST and GREEN WEST will be increased. If not, the scores for these two plans will be decreased. If there is a blue box to the north, the score for BLUE EAST will be increased; otherwise, it will be decreased. If there is at least one green box and one blue box to the north, all of the plans will have their score increased. Similarly, if there are no green or blue boxes to the north, all of the plans will have their scores decreased.

Once a single plan has a higher score than any of the other plans, this is determined to be the best possible plan. If the agent is not able to recognize a plan within a given time frame, the plan-recognition agent then communicates with the observed agent. Whether by plan recognition or communication, the plan-recognition agent begins to follow the recognized plan to aid the observed agent.

Intent recognition differs from plan recognition in that the agent attempts to determine the observed agent's intent instead of the plan it is following. Intent recognition begins similarly to plan recognition. Observations are used to update scores in the intent agent's plan library. In addition to this information, the number of steps that have been observed of a particular plan is recorded. If an action is observed that is not in the plan, the plan score is not reduced. To calculate the score of a plan, the intent agent stores a value called the original score. This is increased every time an action is observed that pertains to the plan. The original score is then combined with the number of plan steps that have been seen to calculate the adjusted score: $\text{adjusted score} = O + (I / (S1 - S2))$, where O is the original score, $S1$ is the number of steps in the plan, $S2$ is the number of steps in the plan that have been observed, and I is an adjustable intent-recognition bonus. When $S1$ and $S2$ are equal, the adjusted score is $O + I$. Future research can change this value to see what the affect in intent recognition will be.

When there is a single plan in the plan library that has a higher adjusted score than any of the other plans, this plan is selected as the most likely to have the same intent as the plan that the observed agent is following. Similar to plan recognition, if a plan is not recognized before the set timeframe, the intent agent communicates with the observed agent to determine its plan. If the intent agent has communicated, it will attempt to aid the observed agent by following the same plan that the observed agent sent in the reply to its query. Intent recognition and plan recognition behave in the same way if communication is necessary. However, if intent recognition was completed prior to the communication, the intent agent will likely execute a different plan than the observed agent to aid it.

An example of this would be if an intent agent is observing an observed agent, which is pushing yellow boxes to the east. The intent agent may decide that the plan it its library with the closest intent is the one which first removes all the obstacles between the yellow boxes and the eastern wall and then proceeds in pushing the boxes.

An agent draws its knowledge about how to interact with the environment from its plan library. Similar to the research from Hongeng and Nevatia [13], we assume that there are two types of events. In our research, they are referred to as simple events and complex events. An agent that is performing plan recognition or intent recognition makes observations in simple events. The simple events used in this simulation are PUSH_BOX_WEST, PUSH_BOX_EAST, PUSH_BOX_NORTH, PUSH_BOX_SOUTH, MOVE_EAST, MOVE_WEST, MOVE_SOUTH, MOVE_NORTH, DESTROY_OBSTACLE, and NONE.

In terms of complex events, based on observations about the environment, the agent then converts these observations into complex actions. For example, an agent is observed to be performing the MOVE_NORTH simple action. If there are red boxes to the north of the agent, this would be interpreted as a MOVE_TOWARDS_RED_BOX complex action. If there were also one or more blue boxes to the north this would also be interpreted as MOVE_TOWARDS_BLUE_BOX. In this way, one observation of a simple event can be translated into one or more complex events.

Plans are collections of complex actions. For example, the plan called PushRedBoxesNorthPlan is made up of the complex actions MOVE_TOWARDS_RED_BOX and PUSH_RED_BOX_NORTH. Another example is the plan called PushRedBoxesEastPlanDF, which consists of MOVE_TOWARDS_RED_BOX and PUSH_RED_BOX_EAST, and DESTROY_DELAY_AHEAD. The DESTROY_DELAY_AHEAD action determines if there are any obstacles between the boxes and the target location, which in this case is the north side of the environment, and then the agent proceeds to remove those obstacles.

One advantage of intent recognition over plan recognition is that the intent-recognition agent's plan library does not have to contain the plan that is being observed for the agent to recognize a plan. Another advantage of intent recognition is that even if the observed plan is in the plan library, the intent-recognition agent is capable of selecting a different plan from its library if the agent determines that both plans have the same intent.

For each data set, the agents performing intent recognition and plan recognition were given the same plan library. Both intent-recognition agents and plan-recognition agents have plans that contain information about destroying obstacles in the environment. However, because the observed agents do not have the capability to destroy obstacles, neither the intent-recognition agents nor the plan-recognition agents will ever make an observation along these lines. However, an intent agent may still choose to perform a plan that includes destroying obstacles if the plan is seen as having the same intent as the observed plan. In this way, intent recognition has an advantage over plan recognition.

If recognition can easily be done, it would be difficult to study the differences between intent recognition and plan recognition. To keep the analysis comparable, all intent recognition and plan-recognition agents have the same plan libraries. These libraries include plans that are identical except for the addition of the DESTROY_OBSTACLE complex action. Observed agents in this research cannot perform this action. This makes the recognition process for plan-recognition agents complex, as illustrated in the following example: for instance, the Plan-recognition agent's Plan Library is $P_0 = \{\text{MOVE_TO_YELLOW_BOX, PUSH_YELLOW_BOX_EAST}\}$ and $P_1 = \{\text{DESTROY_OBSTACLE, MOVE_TO_YELLOW_BOX, PUSH_YELLOW_BOX_EAST}\}$. The plan-recognition agent's observations are the following: T0, observed agent moves towards yellow box; T1, observed agent moves towards yellow box; T2, observed agent pushes yellow box east; T3, observed agent pushes yellow box east; and T4, observed agent pushes yellow box east. Because the order of actions does not matter in our system, the plan-recognition agent does not know whether the observed agent is simply pushing the yellow boxes east and following P_0 or whether the agent is following plan P_1 . The plan-recognition agent is unable to determine whether P_1 is being followed and there are no longer any removable obstacles on the grid or if the observed agent is following the plan steps out of order. In reality, observed agents do not have that capability and the plan-recognition agent is wasting valuable time and resources.

An intent-recognition agent, meanwhile, may select P_0 after the first one or two observations. The agent may also select P_1 after several the first three or four observations. The difference can be affected by factors including on information the agent has collected previously, any observations about the state of the world, and any observations about other agents.

5 Experimental Results

The communication results, shown in Figure 5, indicate that intent recognition is able to communicate fewer times than plan recognition in a multi-agent environment where exact plans cannot easily be determined.

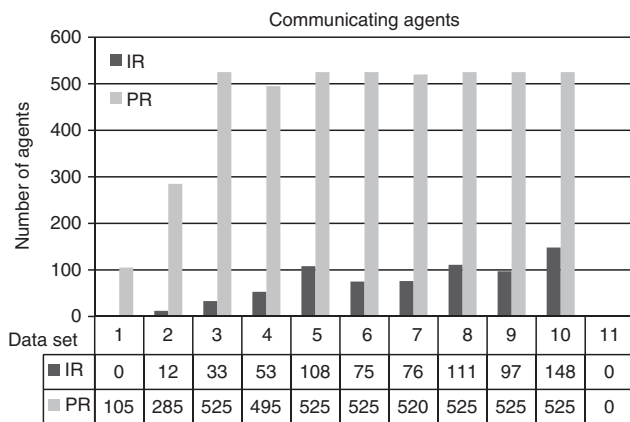


Figure 5. Communicating Agents: Number of Intent-Recognition (IR) Agents and the Number of Plan-Recognition (PR) Agents that Communicated for Each Data Set.

The “no-recognition” group was not included because the observed agents do not initiate communication. Total agents, 5775 agents per recognition type; IR communicating, 713; PR communicating, 4555.

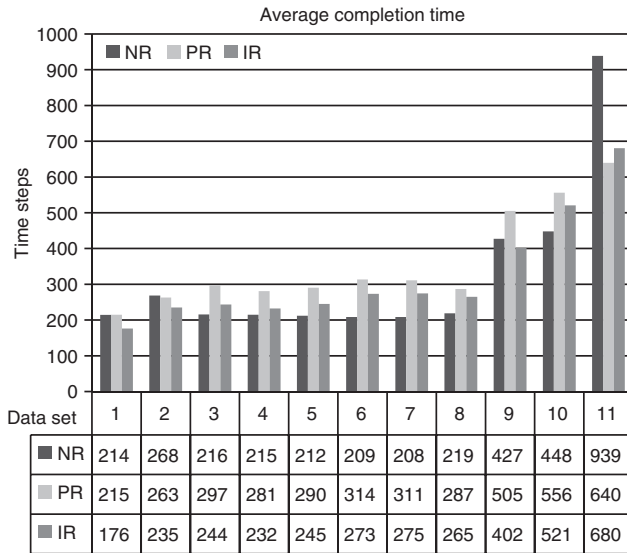


Figure 6. Average Task Completion Time by Data Set and Recognition Type.

The agents performing plan recognition tried to make an exact match between their observations and their plan library. Because of this, often times, the plan-recognition agents were not able to complete the recognition process before reaching the set time limit. Agents performing intent recognition were not looking for an exact match between observations and their plan library. We also found that fewer intent-recognition agents than plan-recognition agents communicated.

The simulation time was recorded as the last time a box was pushed correctly. Simulations where there was no recognition taking place sometime had a faster completion time than the simulations where there was recognition of some sort. This is because in the simulations with no recognition, all agents in the environment began their assigned tasks as soon as the simulation began. In the cases where there were agents performing recognition, these agents would first observe for a period before beginning to execute their tasks. The results are shown in Figures 6 and 7.

A regression analysis was performed using the data from simulation time. The recognition type and the number of agents were used as variables. Using 0 to represent no recognition, 1 to represent plan recognition, and 2 to represent intent recognition, regression was done to see if there was a linear relationship between the recognition types with this dummy variable.

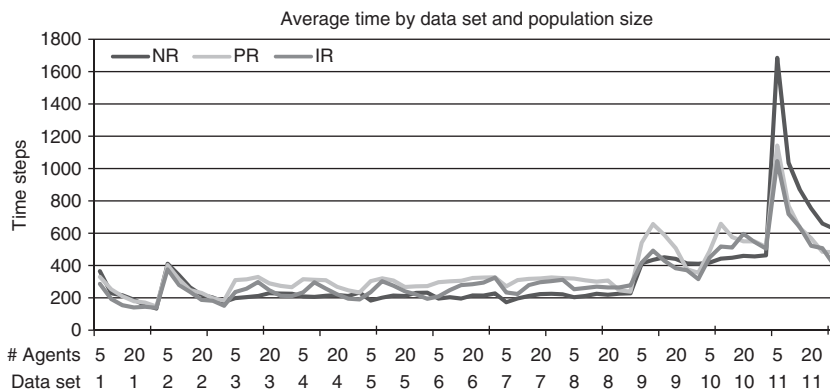


Figure 7. Average Time by Data Set, Population Size, and Recognition Type.

On the horizontal axis, the top number refers to the population size and the lower number refers to the data set number.

The results indicate that there are cases where completion time tends to be higher between groups with recognition and groups without recognition, such as can be seen in data sets 6, 7, 8, and 10. This is because the recognition agents spend the beginning of the simulation observing and do not begin working until they have either determined a plan or communicated. Because of this, recognition agents are usually still working after all of the nonrecognition agents have run out of energy. The goal is to see whether plan recognition and intent recognition are significantly different and which one is correlated with a higher simulation time.

According to our research, in data sets 1 through 9, a significant impact is found by recognition type. Recognition type has a negative impact on the task completion time. This means that intent recognition, with the higher dummy variable value, is associated with a lower completion time for tasks. More details of the time regression are included in the appendix.

The percentage of completion was measured as the number of boxes that were moved to the correct location out of the total number of boxes that were to be moved. The results are shown in Figure 8. The independent variables used were agent population size and regression type. Regression type was a categorical variable with the following values: 0, no recognition; 1, plan recognition, and 2, intent recognition. Data sets 1 and 2 had no variation of percentage of completion. In all data sets aside from 11, there was a significant t-stat for the recognition type. The details of the regression are included in the appendix.

Regression was then done with only the data for the intent recognition and plan-recognition agents. We found that percentage of completion was not statistically different between plan recognition and intent recognition in most data sets. Our research shows that the plan-recognition group and the intent-recognition group had a statistically higher completion percentage than the no recognition group. Although intent-recognition agents did not have a statistically higher completion percentage than intent-recognition agents, they were able to have either a comparable or higher percentage of completion than plan-recognition agents while having much lower number of communicating agents and a faster completion time.

The results indicate that intent-recognition agents performed actions that they had not observed. Intent-recognition agents and plan-recognition agents had additional capabilities when compared with the agents that they were observing. Unlike the observed agents, these agents were able to identify and destroy obstacles in the environment. Because the observed agents did not have the capability to destroy obstacles, the plan-recognition agents did not have any observations along these lines. For this reason, the plan-recognition agents never recognized a plan to destroy an obstacle in it as the correct plan.

Intent agents dynamically chose other actions that they had not observed. Some of these may have been beneficial depending on the overall system goal, whereas others were not. One example is where there were five red boxes and five blue boxes on the grid. In this example, there is one observed agent whose task is to move red boxes north. After the intent agent observes a box being pushed to the north, it may decide that all boxes should be moved to the north. If the observed agent's next task is to push all the boxes to the north,

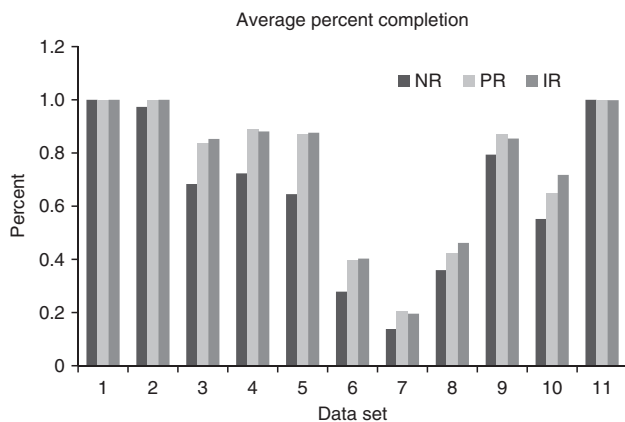


Figure 8. Average Percentage of Completion by Data Set and Recognition Type. NR, group with no recognition; PR, plan recognition group; IR, intent recognition group.

this will help to decrease the task time. However, if the observed agent is actually trying to push the red boxes north and the blue boxes east, this will introduce task error into the system.

Overall, there were 5418 intent-recognition agents that were able to recognize a plan. Of these agents, 5403 of them recognized a plan that had the possibility of assisting the observed agents. 710 of the agents recognized a plan that had the possibility of introducing error into the system. The percentage of completion was comparable between the intent-recognition agents and plan-recognition agents, whereas the intent-recognition agents had faster task completion times and fewer communicating agents.

6 Conclusion

This research defined intent-recognition systems. We designed and built a simulation environment to analyze this concept. We compared intent-recognition systems to plan recognition systems to study the merits of intent recognition. As hypothesized, plan-recognition agents and intent-recognition agents approached the task of aiding fellow agents in different ways. Unlike agents performing plan recognition, intent-recognition agents have the ability to dynamically choose actions to assist other agents in the system using only the knowledge in their plan libraries. Also, intent-recognition systems communicate fewer times compared with plan recognition systems. Although both approaches aim to reduce communication in a system, agents performing intent recognition have the advantage of not trying to make an exact match between observations and the plan library. Thus, communication is further reduced in the case where all agents in the system have differing capabilities and plan libraries.

Costs can be reduced by incorporating intent recognition into multi-agent systems. The frequency of communication is reduced, which in turn reduces the associated costs. Ongoing system tasks can be completed in a shorter amount of time. This allows for more tasks to be completed in any given timeframe. Agents in an intent-recognition system can combine their differing abilities to efficiently complete a wide variety of complex tasks, from box pushing to assembling a space station. The ethical issues of this type of research are important and are beyond the scope of this article. Such issues are discussed in works on robot ethics [20].

Based on this fact that there is substantial advantage in using intent recognition, it is imperative that people who are designing cooperative multi-agent systems consider the intent recognition construct. With a properly seeded plan library, idle agents are able to increase the utility of a system by aiding other agents without human intervention.

This article makes multiple theoretical contributions. Its theoretical contributions are primarily in the field of multi-agent systems.

1. We created and tested a new construct called intent recognition, where agents determine the intent of the agents around them.
2. Agents performing intent recognition have a reduction in the number of communications. This reduces the overall communication in a system.
3. With a properly created plan library, adding agents that perform intent recognition to a system can increase the overall utility, where the utility is domain specific.
4. Intent recognition allows agents to autonomously find ways in which to utilize their idle time. This reduces the amount of time agents are waiting for instructions. It also increases utility of the system by having all agents working at all times.
5. Intent recognition allows agents to assist each other using unobserved actions. Unlike agents performing plan recognition, agents performing intent recognition search for plans in their libraries that have the same goal as the observed actions instead of the same steps. This leads to agents that are able to use actions to achieve the goal that the observed agent may not be capable of.

Bibliography

- [1] A. Agah, Robot teams, human workgroups and animal sociobiology: a review of research on natural and artificial multi-agent autonomous systems, *Adv. Robot. J.* **10** (1996), 523–545.
- [2] D. Avrahami-Zilberbrand and G. A. Kaminka, Hybrid symbolic–probabilistic plan recognizer: initial steps, in: *Proceedings of the AWAIT Workshop on Modeling Others from Observations (MOO-06)*, Boston, MA, 2006.
- [3] D. Avrahami-Zilberbrand, G. A. Kaminka and H. Zorosim, Fast and complete symbolic plan recognition: Allowing for duration, interleaved execution, and lossy Observations, in: *Proceedings of the IJCAI Workshop on Modeling Others from Observations (MOO-05)*, Edinburgh, Scotland, 2005.
- [4] B. Banerjee, L. Kraemer and J. Lyle, Multi-agent plan recognition: formalization and algorithms, in: *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-10)*, 2010.
- [5] W. Briggs and D. Cook, Local planning and teamwork: minimizing communication in multi-agent planning, Technical Report Number TR-CSE-95-001, Department of Computer Science and Engineering, University of Texas at Arlington, Arlington, TX, pp. 688–693, 1995.
- [6] C. Castelfranchi, SILENT AGENTS: from observation to tacit communication, *IBERAMIA-SBIA 2006*, pp. 98–107, Ribeirão Preto, SP, Brazil, 2006.
- [7] Eclipse, 2012. www.eclipse.org.
- [8] A. Gabaldon, Activity recognition with intended actions, in: *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*, pp. 1696–1701, 2009.
- [9] Y. Gal and A. Pfeffer, A language for opponent modeling in repeated games, *AAMAS03 Game theoretic workshop*, pp. 265–272, Melbourne, Australia, 2003.
- [10] C. Geib and S. Harp, Empirical analysis of a probabilistic task tracking algorithm, in: *Proceedings of MOO 2004 Workshop*, pp. 65–71, New York, NY, 2004.
- [11] R. P. Goldman, F. Kabanza and P. Bellefeuille, Plan libraries for plan recognition: do we really know what they model?, in: *Plan, Activity, and Intent Recognition (PAIR) 2010, AAAI Workshop Proceedings*, AAAI, 2010.
- [12] J. R. Hackman, *Groups that work (and those that don't)*, Jossey-Bass, San Francisco, CA, 1990.
- [13] S. Hongeng and R. Nevatia, Multi-agent event recognition, in: *Proceedings of IEEE International Conference on Computer Vision*, pp. 84–91, Vancouver, Canada, 2001.
- [14] Z. Huang, Y. Yang and X. Chen, An approach to plan recognition and retrieval for multi-agent systems, in: *Workshop on Adaptability in Multi-Agent Systems, First RoboCup Australian Open 2003 (AORC-2003)*, Sydney, Australia, 2003.
- [15] M. J. Huber and E. H. Durfee, On acting together: Without communication, in: *Working Notes of the AAAI Spring Symposium on Representing Mental States and Mechanisms*, pp. 60–71, Stanford, CA, 1995.
- [16] M. J. Huber and E. H. Durfee, Deciding when to commit to action during observation-based coordination, in: *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pp. 163–170, AAAI Press, San Francisco, CA, 1995.
- [17] S. S. Intille and A. F. Bobick, A framework for recognizing multi-agent action from visual evidence, in: *AAAI/IAAI'99*, pp. 518–525, Orlando, FL, 1999.
- [18] P. W. Jordan and M. A. Walker, Deciding to remind during collaborative problem solving: Empirical evidence for agent strategies, in: *Proceedings of the 14th Conference of the American Association of Artificial Intelligence*, pp. 16–23, AAAI Press, Portland, OR, 1996.
- [19] C. Y. Lin and J. Y. Hsu, IPARS: Intelligent portable activity recognition system via everyday objects, human movements, and activity duration, in: *Proceeding of AAAI workshop on modeling others from observations*, pp. 44–52, Boston, MA, 1996.
- [20] P. Lin, K. Abney and G. A. Bekey, *Robot ethics: the ethical and social implications of robotics*, MIT Press, Cambridge, MA, 2011.
- [21] A. Majchrzak and L. Gasser, Toward a conceptual framework for specifying manufacturing workgroups congruent with technological change, *Int. J. Comput. Integr. Manuf.* **5** (1992), 118–131.
- [22] D. Mammen and V. Lesser, A test bed for the evaluation of multiagent communication and problem-solving strategies, in: *Constraints and Agents: Papers from the 1997 AAAI Workshop*, Menlo Park, CA, AAAI Press, Technical Report WS-97-05, pp. 32–38, 1997.
- [23] E. Marhasev, M. Hadad and G. A. Kaminka, Non-stationary hidden semi Markov models in activity recognition, in: *Proceedings of the AAAI Workshop on Modeling Others from Observations Boston, MA (MOO-06)*, pp. 53–60, 2006.
- [24] M. J. Mataric, Issues and approaches in the design of collective autonomous agents, *J. Robot. Auton. Syst.* **16** (1995), 321–331.
- [25] B. McEleney and G. O'Hare, Planning to replan in a multi-agent environment, in: *Proceedings of the 23rd Workshop of the UK Planning and Scheduling Special Interest Group* December 20–21, 2004, Cork, Ireland, 2004.
- [26] B. McEleney and G. O'Hare, A planner for plan construction plans for multi-agent plans, in: *Proceedings of 15th Irish Conference on Artificial Intelligence and Cognitive Science (AICS'04)* September 8–10, 2004, Castlebar, Ireland, 2004.
- [27] B. McEleney and G. O'Hare, Efficient dialogue using a probabilistic nested user model, in: *Proceedings of the Fourth IJCAI Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, August 1, 2005, Edinburgh, Scotland, 2005.

- [28] M. Molineaux, D. Aha and G. Sukthankar, Beating the defense: using plan recognition to inform learning agents, in: *Proceedings of 22nd Florida Artificial Intelligence Research Society (FLAIRS)*, pp. 257–262, 2009.
- [29] Repast, 2012. <http://repast.sourceforge.net>.
- [30] M. O. Riedl and R. Young, An intent-driven planner for multi-agent story generation, in: *AAMAS '04 Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, Vol. 1, pp. 186–193, 2004.
- [31] I. Rish, An empirical study of the naive Bayes classifier, in: *Proceedings of IJCAI-01 workshop on Empirical Methods in AI, International Joint Conference on Artificial Intelligence 2001*, pp. 41–46, 2001.
- [32] C. Tsatsoulis, N. Ahmad, E. Komp and C. Redford, Using a contract net to dynamically configure sensor webs, *IEEE Aerospace Conference*, Big Sky, MT, pp. 1–6, 2008.
- [33] M. A. Walker and O. Rambow, The role of cognitive modeling in achieving communicative intentions, in: *Proceedings of The 7th International Conference on Natural Language Generation*, New Forest, England, pp. 171–180, 1994.
- [34] M. Wooldridge, *An introduction to multiagent systems*, John Wiley & Sons, Chichester, UK, 2002.
- [35] P. Xuan, V. Lesser and S. Zilberstein, Formal modeling of communication decisions in cooperative multi-agent systems, in: *Proceedings of the Second Workshop on Game Theoretic and Decision Theoretic Agents (GTDT 2000)*, Boston, MA, 2000.
- [36] P. Xuan, V. Lesser and S. Zilberstein, Communication decisions in multi-agent cooperation: model and experiments, in: *Proceedings of the Fifth International Conference on Autonomous Agents*, Montreal, Canada, pp. 616–623, 2001.
- [37] G. Ya'akov, S. Reddy, S. Shieber, A. Rubin and B. Grosz, Plan recognition in exploratory domains, *Artif. Intell.* **176** (2012), 2270–2290.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.