Collective Intention Recognition and Elder Care

Han The Anh and **Luís Moniz Pereira**
Centro de Inteligência Artificial (CENTRIA)
Departamento de Informática, Faculdade de Ciências e Tecnologia
Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
$h.anh@fct.unl.pt, lmp@di.fct.unl.pt$

Abstract

The contribution of this paper is twofold. First, we present a new method for collective intention recognition based on mainstream philosophical accounts. Second, we extend our previous Elder Care system with collective intention recognition ability for assisting a couple of elderly people. The previous system was just capable of individual intention recognition, and so it has now been enabled to deal with situations where the elders intend to do things together.

Introduction

In the last twenty years there has been a significant increase of the average age of the population in most western countries and the number of elderly people has been and will be constantly growing. For this reason there has been a strong development of supportive technologies for elderly people living independently in their own homes, for example, RoboCare Project (Cesta and Pecora 2004) – a project developing robots for assisting elderly people's living, SINDI – a logic-based home monitoring system (Mileo, Merico, and Bisiani. 2008) and PHATT – a framework developed for addressing a number of desired features for Elder Care domain (Geib 2002).

For the Elder Care application domain, in order to proactively provide contextually appropriate help for elders, it is required that the assisting system have the ability to observe the actions of the elders, recognize their intentions, and then provide suggestions on how to achieve the recognized intentions on the basis of the conceived plans. In (Pereira and Han 2009a; 2010), we have presented a system focusing on the latter two steps in order to design and implement an Elder Care logic programming based assisting system. The first step of perceiving elders' actions is taken for granted. For elders' intention recognition based on their observable actions, we employ our work on Intention Recognition (IR) using Causal Bayes Networks and plan generation techniques, described in (Pereira and Han 2009c). The IR component is indispensable for living-alone elders, in order to proactively provide them with timely suggestions.

However, since this system is only capable of individual IR, it is unable to deal with the problem domain where a couple of elderly people live alone in their apartment. In this domain, there are cases where the elders intend to do things together, i.e. having a collective intention, and it is likely that individual intentions do not make sense or provide useful information. As most researchers in philosophy (Tuomela and Miller 1988; Searle 1990; Bratman 1992) and multi-agent systems (Kanno 2003) agree, collective intentions (or joint intentions; we-intentions; shared intentions) are not summative. A collective intention of a group of agents cannot be reduced to a mere summation of the individual intentions of the agents. It involves a sense of acting together and willing something cooperatively, thus some kind of "glue", e.g. mutual beliefs or mutual expectations, must exist amongst the agents.

We will present a new method for collective IR and extend our previous system with this ability to take care of the situation where there is a couple of elderly people staying alone in their apartment [1]. In order to assist the couple properly, it is important to have in such cases the ability to detect whether they have some collective intention (and recognize it); otherwise, individual IR should be performed. It is important to stress that individual IR should not be performed unless collective intentionality is confirmed not to exist.

Collective Intention and Recognition

Collective intention is one of the active research issues generally discussed in philosophical and multi-agent system literature. Most researchers agree that collective intentions are not summative, i.e. cannot be reduced to a mere summation of individual intentions (Bratman 1992; Tuomela 2005; Searle 1990). Collective intentions involve a sense of acting and willing something together cooperatively. There must be some kind of "glue" supplementing the separate individual intentions in order for agents to partake in a collective intention, e.g. mutual beliefs according to Tuomela and mutual awarenesses according to Searle. In (Tuomela and Miller 1988; Tuomela 2005), the collective intention (or we-intention as he used) of a group of agents is defined as individual intentions of the agents plus their mutual beliefs. Briefly, agent A and B intend to do some task X cooperatively if the following "glue" conditions for A (and the sym-

[1] There may be more than two elders living together, but the scenario of an elderly couple requiring care service is more usual. Furthermore, we shall see that our system can be naturally extended to the general case of multiple elderly users.

metrical conditions for B) hold

(a) A intends to do his part of X
(b) A believes that B will do his part of X
(c) A believes that B believes that he will do his part of X

Following Tuomela, Kanno et al. presented a bottom-up approach to collective IR (Kanno 2003). To recognize the collective intention of a group of agents, the individual intentions and beliefs of the constituents are inferred first. Then, the collective intention is inferred by checking for consistencies amongst those inferred mental components.

The main disadvantage of this bottom-up approach is that it is confronted with a combinatorial problem of possible combinations of individual intentions and beliefs to form collective intentions. Given the situation at hand, each agent may have several conceivable intentions and beliefs, but there are not many combinations of them forming conceivable collective intentions.

To tackle this issue, we propose a top-down approach to collective IR. The recognition process starts by inferring the possible collective intentions assuming that they were had by a virtual plural agent representing the group and abreast of all group activity. Then, we figure out which of them is a genuine one by checking whether there is any activity "glue" information linking the agents' individual intentions or activities. The above assumption is inspired and validated by Searle's account of collective intention (Searle 1995; 1990). According to him, collective intentionality is non-summative, but remains individualistic. With the presupposition of mutual awarenesses—namely that each agent supposes the others are like himself and that they have similar awareness of him as an agent like themselves—it allows for the possibility of a single plural agent or "brain in a vat" having the collective intention of the group. Thus, if a group of agents had a collective intention, this intention could be recognized as if it was had by a single agent. For this we can use any existing individual IR methods.

Now let us look at the second step of confirming which of the inferred collective intentions is the genuine one. For intention recognition's sake, what we are interested in (and actually all what we can have) are the actions or their effects in the environment resulting from the "glue" mental attitudes (mutual beliefs or mutual awarenesses) amongst the agents. An intermediate stage between having such mental attitudes and actual activity is that the agents form some mutual expectations between each other which reflect their attitudes. Thus, if and when having a collective intention, each agent in the group should act according to his expectations to other constituents. Namely, when working together towards achieving a collective task (intention), an agent may expect from another agent (or a group of other agents) who is responsible for producing some result for his input. From the opposite side, the result-producer agents expect their result-consumer agents to use the result as expected. In addition, if some agents are doing the same task (no one needs input result from other), then they expect from each other to commit to doing that task. If an agent wants to do something else (e.g. have a break), others would expect him to tell them about that. Otherwise, they would complain.

In this work we assume that a priori domain knowledge is specified in the form of a library containing the set of possible plans and expectation actions. Automatically learning these models from data is beyond its scope.

Based on the above discussion, next we show in detail a top-down method for collective IR, employing the individual IR approach in (Pereira and Han 2009c) for illustration.

Method for Collective Intention Recognition

The method for recognizing collective intentions consists of two steps:

1. From the observations (actions or their effects in the environment of all agents in the group) infer the intentions as if these observations came from a plural agent; then
2. Figure out which of the recognized intentions is a genuine collective intention by checking if there are actions reflecting the mutual expectations between the agents, ignoring irrelevant actions for each considered intention.

The first step of recognizing intention of a single agent is realized using our previous work on IR via Causal Bayes Nets (CBN) and a plan generator (or a given plan library) (Pereira and Han 2009c). The CBN is used to generate conceivable intentions of the intending agent and compute their likelihood conditional on the initially available observations, and so allow to filter out the much less likely ones. The plan generator thus only needs to deal with the remaining more relevant intentions, because more probable or credible, rather than all conceivable intentions. In the sequel the network structure for IR is recalled. We assume the readers are familiar with the basic concepts of CBNs, which can be obtained from (Pearl 2000; Pereira and Han 2010) .

Network Structure for Intention Recognition

The first phase of the IR system is to find out how likely each conceivable intention is, based on current observations such as observed actions of the intending agent or the effects its actions had in the environment. A conceivable intention is the one having causal relations to all current observations. It is brought out by using a CBN with nodes standing for binary random variables that represent causes, intentions, actions and effects.

Intentions are represented by those nodes whose ancestor nodes stand for causes that give rise to intentions. Intuitively, we extend Heinze's tri-level model (Heinze 2003; Pereira and Han 2010) with a so-called pre-intentional level that describes the causes of intentions, used to estimate prior probabilities of the intentions. However, if these prior probabilities can be specified without considering the causes, intentions are represented by top nodes (i.e. nodes that have no parents). These reflect the problem context or the intending agent's mental state.

Observed actions are represented as children of the intentions that causally affect them. Observable effects are represented as bottom nodes (having no children). They can be children of observed action nodes, of intention nodes, or of some unobserved actions that might cause the observable effects that are added as children of the intention nodes.

The causal relations among nodes of the CBNs (e.g. which causes give rise to an intention, which intentions trigger an action, which actions have an effect), as well as their Conditional Probability Distribution (CPD) tables and the distribution of the top nodes, are specified by domain experts. However, they might be learnt mechanically.

Example 1 (Elder Care) *A couple of elderly people stay alone in their apartment. The IR system observes that they are both now in the kitchen. The man is looking for something and the woman is holding a kettle. In order to assist them, the system needs to figure out what they intend to do, whether cooperatively or individually. The possible collective intentions are: making a drink or cooking. The CBN with CPD tables is provided in Figure 1.*

In our work, the probabilistic inference in CBNs is automatically done with P-log – a probabilistic logic programming system (Baral, Gelfond, and Rushton 2009; Han, Ramli, and Damásio 2008). The probabilities that the elders have the collective intentions of *cooking (cook)* and *making a drink (mD)* given the observations that the man is looking for something and the woman is holding a kettle, are computed with the following P-log queries, respectively:

$$?-pr(i(cook,t) \mid (obs(look(t)\ \&\ obs(holdKettle)),\ V_1).$$
$$?-pr(i(mD,t) \mid (obs(look(t)\ \&\ obs(holdKettle)),\ V_2).$$

The result is: $V_1 = 0.478$; $V_2 = 0.667$. It means that the collective intention of making a drink is more likely and should be examined first. However, it is still necessary to look at the other collective intention since it is not much less likely. In this case, we used the general CBN default for the problem domain. In a given situation, more information may be available, and we should be able to render the CBN more specific to the situation.

Situation-sensitive CBNs Undoubtedly, CBNs should be situation-sensitive since using a general CBN for all specific situations of a problem domain is unrealistic and most likely imprecise. For example, different elders will have different conditions and habits that need to be taken into account to recognize their intentions. Also, place, time of day, temperature, etc. need to be considered. However, consulting the domain expert to manually change the CBN w.r.t. each situation is also very costly or unfeasible.

In (Pereira and Han 2009c) is provided a way to construct situation-sensitive CBNs, i.e. ones that change according to the given situation. It uses Logic Programming techniques to compute situation specific probabilistic information that is then updated into a CBN—which is general for the problem domain. The CBNs themselves are also encoded with P-log (Baral, Gelfond, and Rushton 2009; Han, Ramli, and Damásio 2008), which supports coherent updates. That is the main reason why P-log is used rather than the standard graphical model inference (besides its efficient implementation for multiple probabilistic querying).

Example 2 (Elder Care (cont'd)) *In the scenario provided in the previous example, the CBN may vary depending on some observed factors, for example, the time of day, of the elders' last drink or last meal, etc. We design a logical component for the CBN to deal with those factors:*

```
pa_rule(pa(hg(t),d_(0,1)),[])
      :-time(T), eat(T1), T-T1 < 1.
pa_rule(pa(hg(t),d_(9,10)),[])
      :-time(T), last_eating(T1), T-T1 > 3.
pa_rule(pa(thty(t),d_(1,10)),[])
      :-time(T), drink(T1), T1-T < 1.
pa_rule(pa(thty(t),d_(9,10)),[])
      :-time(T), last_drink(T1), T1-T > 3.
```

Basically, probabilistic information is given by $pa/2$ rules. For example, the rule $(pa(hg(t), d\_(9, 10)) \leftarrow)$ means that the probability of being hungry (i.e. $hg(t)$) is $9/10$ unconditionally (since the rule body is empty). We provide a reserved $pa\_rule/2$ predicate which takes the head and body of some $pa/2$ rule as its first and second arguments, respectively, and includes preconditions for its activation in its own body. Now, a situation is given by asserted facts representing it and, in order to find the probabilistic information specific to the given situation, we simply use the XSB Prolog built-in *findall/3* predicate to find all true $pa/2$ literals expressed by the $pa\_rule/2$ rules with true bodies in the situation.

For example, suppose that the current time is 18 ($time(18)$) and the last time the elders ate was half an hour before ($last\_eating(17.5)$). But they did not have any drink for 3 hours (e.g. $last\_drink(14)$). Those three facts are asserted. Hence, the following two $pa\_rule/2$ literals are true, and are updated into the general CBN

```
pa_rule(pa(hg(t),d_(0,1)),[]).
pa_rule(pa(thty(t),d_(9,10)),[]).
```

Now the result is: $V_1 = 0.0206$; $V_2 = 0.9993$. This time the single collective intention of making a drink should be sought for confirmation in the next stage, since the one of cooking is very unlikely.

Confirming Collective Intention

The next step is to confirm whether the recognized intention is actually a collective intention of the group of agents. This is done by checking if there are expectation actions between the agents which reflect their mutual expectations.

Let $\{a_1, ..., a_n\}$ be the set of agents and A the plural agent representing these agents (i.e. having all their actions). Suppose W is an intention of A, recognized from the previous step, and L is the list of plans achieving W.

Let $P = [p_1, ..., p_k]$ be a plan in L. We assume here for simplicity that plans are sequential. As we shall see later, the agents doing the same task are grouped. If there are agents in the group doing concurrent actions, the plural agent is assumed to do the work of one agent before the other.

Now we can determine the assigned subplan of each agent towards achieving the collective intention of the whole group, by looking at each agent's actions.

Let $s_i, 1 \leq i \leq n$, be the first action of agent a_i in P. Determine indices $d_i, 1 \leq i \leq n$, such that $p_{d_i} = s_i$. Group the agents having the same first action, i.e. with the same index d_i. They are doing the same task or at least some part together. Suppose we obtain m groups $g_1, ..., g_m$, and g_t is responsible for the subplan $[p_{j_t+1}, ..., p_{j_{t+1}}]$ where $1 \leq t \leq m$ and $0 = j_1 < j_2 < ... < j_{m+1} = k$. It is easily seen that

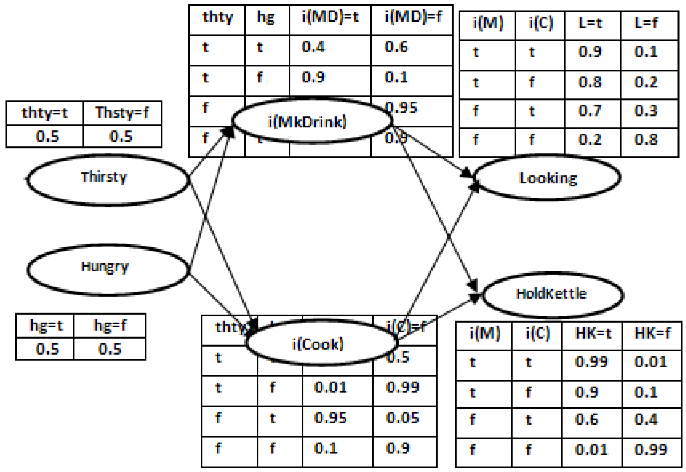

Figure 1: Elders Collective intentions CBN

the grouping is unique for a given set of agents and a given plan.

In order to check if there are expectation actions that reflect agents' mutual expectations, we consider two cases: mutual expectations between agents in a group and between agents in consecutive or subsequent subplans groups [2]. This way, the number of interactions amongst the agents that need to be observed (e.g. by an activity recognition system) is considerably reduced. Furthermore, the number of possible expectation actions between two particular agents is reduced and being more specified.

Expectations inside Group When some agents are assigned or intended to do something together, they expect from each other to do the task. Thus, the expectation actions can be a "complain" action when a group member "deviates" from the task without an "inform" action.

An example where two people have an intention to walk together (Tomasello 2008), if one of them changes his/her direction without informing the other, he/she obviously will be complained on by the other. This is to differentiate from the situation where those two people walk in step by chance.

In short, we say a group of agents having the same subplan (or part of it) actually work cooperatively if we observe expectation action "complain" if there is any agent deviating from the assigned task.

When agents cooperatively work together, if some agent cannot manage his task, we usually observe some kind of "help" action. However, we believe that kind of action is mostly preceded by a "complain" action. Inability to do the assigned task is also one kind of deviation.

The term "deviation" is used here as a general term, and should be specified for concrete application domains. For example, in spatial domain (Sukthankar 2007; Devaney and Ram 1998) an agent is said to deviate if he is not in his assigned position or does not move in the right direction or does so with wrong velocity.

Expectations between Consecutive Groups When agents are in consecutive groups, says g_t and g_{t+1}, responsible for the subplans $[p_{j_t+1}, ..., p_{j_{t+1}}]$ and $[p_{j_{t+1}+1}, ..., p_{j_{t+2}}]$, respectively, agents from the first would expect some result from the second, who in turn expect the agents from the first to use their result. If the agents from one of the groups "deviated" from their task or did not finish it as assigned, the agents from the other group would "complain". For example, two people want to make coffee, one is assigned to boil water and other is assigned to look for the coffee. Having boiled the water, the first may have the "expect result" action "ask for the coffee", or he may have "complain" action if the second person could not find the coffee after a while. Some "help" action may follow, but we could conclude the collective intentionality solely with the "complain" action. However, if the second found the coffee earlier, he has an "expect use" action, e.g. "this is the coffee, use it!".

Let $result_t$ be the assigned result of group g_t. Usually, this result comes from the last action, $p_{j_{t+1}}$, of the group subplan. For IR, we assume that the set of possible actions reflecting the expectation of $result_t$ of the action $p_{j_{t+1}}$, denoted by $expect\_result_t$, and the set of possible actions reflecting the expectation of using $result_t$, denoted by $expect\_use_t$, are given.

Then, we say these two groups of agents are working towards achieving a collective intention if either agents in g_{t+1} have some action belonging to $expect\_result_t$, or agents in g_t have some action belonging to $expect\_use_t$, or there are "complain" actions from one of the groups. As long as the collective intentionality inside each group is confirmed, a single expectation action observed between the two groups is enough to conclude that they have a collective intention. Usually, it is useful to identify one of the result-delivery and

[2]It might be the case that one group interacts (produces or consumes results) with more than one other group, but that is not considered in this paper.

result-receiver agents.

Elder Care with Intention Recognition

We combine individual and collective IR to design an assistive system for the domain where a couple of elderly people live alone in their apartment. However, note that the presented collective IR method is applied for the general case of an arbitrary set of agents, the assistive system should be naturally generalized. In order to provide appropriate assistance, the system should be able to find out individual intentions when observing only individual activity (e.g. when the other is absent from the apartment) as well as detect whether there is collective intentionality when observing both elders' activity. In the latter case, if there is no collective intentionality detected, the system should perform individual IR.

When having recognized the intention of the elders, whether individual or collective, the Evolution Prospection system described in (Pereira and Han 2009b; 2010) can be employed to provide appropriate suggestions to achieve the recognized intention, taking into account elders' preferences, health reports, future scheduled events, etc. However, that will not be discussed in this paper.

We continue the previous example for illustration.

Example 3 (Confirming Collective Intention) *Suppose "making a drink" was the collective intention recognized from prior step. We check if it is a genuine one.*

Let us consider a simple plan achieving that intention *[take the kettle, fill it up with water, boil the water, look for tea or coffee, put it into the boiled water]*. Hence, the woman's subplan is *[take the kettle, fill it up with water, boil the water]* and the man's is *[look for tea or coffee, put it into the boiled water]*. The assigned result of the woman is to provide some *boiled water*. The man's expectation is that of *boiled water* from the woman, thus may have some "expect result" actions, e.g. ask whether the water is ready or get the water from the woman. Or, if after a while the woman could not boil the water or she was doing something else, the man would complain. If such "expect result" or "complain" action occurs, we can conclude that they really have a collective intention of making a drink. Otherwise, e.g., the man does not show any expectation for the water that the woman has boiled, then we can conclude that that is not their genuine collective intention, even if later he might use the boiled water for his own purpose. We emphasize that there are necessary actions showing the mutual expectations of results and usage of results when the agents have a collective intention towards achieving some task. Our system also allows for expectations to be updated as a result of changing circumstances, new observations or, say, memory loss, or even of desisting from a common intention. Consequently, expectations and counter-expectations can evolve, subject to preferences and degrees of commitment.

Now suppose that the system has found out that there is no collective intention amongst the elders, and the man keeps looking for something. To assist him, the system should then figure out what is his individual intention.

Example 4 (Man's intentions) *Suppose the possibilities are: book, drink, remote control, cook, make drink.*

The readers are referred to our previous work in (Pereira and Han 2010) for a similar example. We only want to make a note that the information obtained from the collective IR process should be updated into the individual IR process. For example, if the woman made the drink for both, then the intention of "make a drink" and "look for a drink" should be excluded.

Remarks on Complexity

The first step of the collective IR system is that of from the observations to infer the intentions as if these observations came from a plural agent. This step can be done by any existing individual IR systems, hence we do not evaluate its complexity here.

In the second step of confirming collective intention, we suggested a grouping method based on the plan achieving the recognized intention. That method enables to reduce the number of interactions needed to be observed as well as to better focus on smaller groups of agents, with smaller sets of possible expectation actions.

In fact, for a set of n agents, the number of interactions to be observed is $\frac{n(n-1)}{2}$. Applying the grouping method, suppose we obtain m groups of n_j $(1 \le n_j \le n,\ 1 \le j \le m)$ agents, where $\sum_{j=1}^{j=m} n_j = n$. In this case, the number of interactions to be observed is

$$(\sum_{j=1}^{m} \frac{n_j(n_j-1)}{2}) + m - 1 = \frac{n(n-1)}{2} - S$$

where

$$S = (\sum_{1 \le j < k \le m} n_j n_k) - m + 1$$

is the number of interactions not needed to be observed, compared with the case without grouping. When the groups are equally divided, we reduce approximately m times the number of interactions to be observed.

Furthermore, the grouping divides the big set of agents into smaller groups, which we believe will enable them to be more easily observed (e.g. by an activity recognition system). Also, this approach allows specifying which are the expectation actions that need to be recognized between a particular pair of agents. Are they in the same group? Are they in consecutive groups? Or else? In the initial set of agents (without grouping), a bigger set of possible expectation actions needs to be put under consideration for any pair.

Related Works

We are not the first to suggest the use of the plural agent concept for collective intention or plan recognition. Indeed, most of the works in multi-agent plan recognition rely on the assumption that the plan is carried out by a single entity (i.e. a plural agent) such as a team, a group, a troop, and so on, and use the same recognition methods of an individual agent; e.g. in (Devaney and Ram 1998) and (Sukthankar and Sycara 2008), just to name a few. However, to the best of our knowledge, none of these works has addressed the necessary cognitive underpinnings amongst the constituents—such as mutual beliefs or mutual awarenesses—in order to confirm the existence of a collective intention, but just a coincidentally formed one instead. That mutual confirmation has actually been the main concern of the philosophical community regarding collective intentionality.

As a consequence, the work in multi-agent plan recognition has been restricted to considering only sets of agents with an initially assigned collective intention, such as football teams or army troops. They could recognize the collective intention of an arbitrary set of agents assuming that it existed; but they can not figure if it actually did so be designed. For the Elder Care domain concerning multiple users, there are often the cases where the elders' actions accidentally form a plausible plan—which achieves a conceivable intention—but actually each of them is following his/her own intention. Thus, the extant multi-agent plan recognition methods are not appropriate to tackle this issue.

The collective IR method we have presented—namely the part of confirming whether a given intention is a genuine collective intention of a given group of agents—supplements the existing multi-agent intention recognition methods to deal with arbitrary sets of agents, without an initially assigned collective task or intention.

Conclusion and Future Work

We have shown a top-down approach to collective intention recognition, which starts with the assumption that the intention is had by a plural agent that has all the activity of the group of agents being considered. Then, that intention is inferred using an individual intention recognition system. The inferred intention undergoes a confirmation process that checks whether there are actions achieving the intentions and reflecting mutual expectations amongst the agents formed by having a collective intention. This top-down approach to collective intention overcomes the combinatoric issue confronted by bottom-up approaches.

To that effect, we have extended our previous Elder Care assistive system with this ability of collective intention recognition in order to deal with the problem domain where there is a couple of elderly people living alone, assisted by a helping intention recognition system.

Although we have shown the applicability of our collective intention recognition method to an extended example in the Elder Care domain, there remains to apply it to other more complex domains where teamwork occurs, e.g. security, social networks and sport settings.

References

Baral, C.; Gelfond, M.; and Rushton, N. 2009. Probabilistic reasoning with answer sets. *Theory and Practice of Logic Programming* 9(1):57–144.

Bratman, M. E. 1992. Shared cooperative activity. *The Philosophical Review* 101(2):327–341.

Cesta, A., and Pecora, F. 2004. The robocare project: Intelligent systems for elder care. In *AAAI Fall Symposium on Caring Machines: AI in Elder Care*.

Devaney, M., and Ram, A. 1998. Needles in a haystack: plan recognition in large spatial domains involving multiple agents. In *Procs. 15th national conference on Artificial intelligence*. AAAI.

Geib, C. W. 2002. Problems with intent recognition for elder care. In *Procs. AAAI Workshop Automation as Caregiver*.

Han, T. A.; Ramli, C. K.; and Damásio, C. V. 2008. An implementation of extended p-log using xasp. In *Procs. Intl. Conf. Logic Programming (ICLP08)*, 739–743. Springer LNCS 5366.

Heinze, C. 2003. *Modeling Intention Recognition for Intelligent Agent Systems*. Ph.D. Dissertation, The University of Melbourne, Australia.

Kanno, T., N. K. F. K. 2003. A method for team intention inference. *International Journal of Human–Computer Studies* 33(4):393–413.

Mileo, A.; Merico, D.; and Bisiani., R. 2008. A logic programming approach to home monitoring for risk prevention in assisted living. In *Procs. Intl. Conf. Logic Programming (ICLP08)*, 145–159. Springer LNCS 5366.

Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge U.P.

Pereira, L. M., and Han, T. A. 2009a. Elder care via intention recognition and evolution prospection. In *18th Intl. Conf. on Applications of Declarative Programming and Knowledge Management (INAP'09)*.

Pereira, L. M., and Han, T. A. 2009b. Evolution prospection in decision making. *Intelligent Decision Technologies* 3(3):157–171.

Pereira, L. M., and Han, T. A. 2009c. Intention recognition via causal bayes networks plus plan generation. In *Progress in Artificial Intelligence, Procs. 14th Portuguese Intl. Conf. on Artificial Intelligence (EPIA'09)*, 138–149. Springer LNAI 5816.

Pereira, L. M., and Han, T. A. 2010. Intention recognition with evolution prospection and causal bayesian networks. In *Computational Intelligence for Engineering Systems 3: Emergent Applications*. Springer. (forthcoming).

Searle, J. R. 1990. Collective intentions and actions. In *Intentions in Communication*. 401–616.

Searle, J. R. 1995. *The Construction of Social Reality*. New York: The Free Press.

Sukthankar, G., and Sycara, K. 2008. Robust and efficient plan recognition for dynamic multi-agent teams. In *Procs. Intl. Conf. on Autonomous Agents and Multiagent Systems (AAMAS'08)*, 1383–1388.

Sukthankar, G. R. 2007. *Activity Recognition for Agent Teams*. Ph.D. Dissertation, Robotics Institute School of Computer Science, Carnegie Mellon University.

Tomasello, M. 2008. *Origins of Human Communication*. MIT Press.

Tuomela, R., and Miller, K. 1988. We-intentions. *Philosophical Studies* 53:367–390.

Tuomela, R. 2005. We-intentions revisited. *Philosophical Studies* 125(3):327–369.