Probabilistic Plan Inference for Group Behavior Prediction

Wenji Mao[1] Jonathan Gratch[2] Xiaochen Li[1]

[1]*Key Laboratory of Complex Systems and Intelligence Science, Institute of Automation, Chinese Academy of Sciences*

[2]*Institute for Creative Technologies, University of Southern California*

E-mail: {wenji.mao@ia.ac.cn; gratch@ict.usc.edu; xiaochen.li@ia.ac.cn}

ABSTRACT In this paper, we present a decision-theoretic approach to plan inference. Based on the assumption that a rational agent will adopt a plan that maximizes its expected utility, we view plan inference as reasoning about the decision-making strategy of the observed agent. Different from the previous related work, our approach explicitly takes the observed agent's preferences into consideration, and computes the expected utilities of plans to disambiguate competing hypotheses. We use online group data to construct the domain plan library and empirically evaluate our approach in group behavior prediction. The experimental results show the effectiveness of our approach in inferring intentions and goals of entities.

Keywords *Plan Representation*, *Probabilistic Reasoning*, *Intention Recognition*, *Group Behavior Prediction*

1 INTRODUCTION

A growing number of applications have sought to incorporate automatic reasoning techniques into intelligent agents. The reasoning ability is critical for an intelligent agent to make sense of the world it situates and the behavior of others. With the advance of multi-agent interactive environments, user-aware adaptive interfaces and social computing systems that involve people, organizations and rich social interactions, it is increasingly important to model and reason about the behavior of other entities. In general, by inferring the behavior of other entities, causal reasoning helps understand and explain the behavior of others and thus can facilitate various forms of social interactions. It can also facilitate behavior representation and modeling, simulating agent society, and predicting future behavior.

In constructing the inferential mechanism for intelligent agents, we adopt a plan-based approach. Plan representations are typically used by many intelligent systems, especially agent-based systems. Plans provide a concise description of the causal relationship between goals, events and states. They also provide a clear structure for exploring alternative courses of actions, and interactions between future activities. Such representations have several key advantages: recognizing the relevance of events to agents' goals and plans – key for intention recognition; assessing agents' freedom and choice in acting – key for the assessment of power and control; and detecting how an agent's plan facilitates or prevents the plan execution of other agents – key for the detection of plan interventions.

Using plan representations, in this paper, we present a probabilistic reasoning approach to *intention recognition* (or *plan recognition*). Plan recognition is the process of inferring the plans and goals of the observed agent based on a sequence of observations, usually with the help of a set of predefined *recipes* (called *plan library*), which comprises the knowledge and action steps that could be performed by the agent. Previous research on plan recognition has largely ignored the influence of world states (especially the observed agents' preferences over states) on the recognition task, and the proposed models were rarely tested in real scenarios. In our approach, we view plan recognition as inferring the decision-making strategy of the observed agent (or a group of agents) and *explicitly* take the observed agents' preferences into consideration in the recognition process.

To test the effectiveness of our approach, we apply our approach to the prediction of group behavior. *Group behavior prediction* is an emergent research and application field, which studies computational methods for the automated prediction of what a group might do. As many applications could benefit from forecasting an entity's behavior for decision making, assessment and training, it is gaining increasing attention in recent years. Its applications range from homeland and national security to government policy evaluation and market analysis to pandemic and disaster response planning. Research on group behavior prediction has focused on building predictive models for socio-cultural-political modeling, which includes three key issues: data collection, model construction and forecasting using the model [1].

In group behavior prediction, the predictive models were typically constructed using machine learning methods. Martinez *et al.* [1] propose a new algorithm for model construction, CONVEX, which is more computationally efficient than the previous algorithms. CONVEX is essentially a variant of *k*NN method. Although machine learning can automate the construction of predictive models, the structured data used by the learning algorithms need to be manually coded by humans. For example, CONVEX is based on the MAROB datasets, a collection of historical behavioral records of ethnopolitical groups as well as environmental factors associated with the group behavior. The datasets were handcrafted by domain experts through examining large volumes of trusted news reports [1]. This process is not only painstaking but prohibitive for massive data collection.

Group behavior prediction provides an ideal testbed for practicing and evaluating plan inference approaches. There are huge amounts of group data available online. Recent progress has made it possible to automatically extract plan knowledge (i.e., actions, their preconditions and effects) from online raw textual data and construct group plans by means of planning algorithm, albeit in the restrictive security informatics domain [2]. Compared with machine learning-based methods, plan-based inference provides additional advantages for representing, analyzing and explaining behavior prediction results. With the structural plan representation, plan-based inference can not only come up with prediction results, it can also inform the goals and intentions behind the predicted behavior. Plans are more expressive in representing behavioral patterns, group strategies and alternative courses of actions, which provide important information for behavior analysis. Thus in contrast to data-driven approaches, plan inference can better explain the prediction models and results, and improve interpretability.

In addition to proposing a model, we conduct experiment to evaluate our work in group behavior prediction. Based on the realistic group data, we construct plan library and compare our model predictions with the results generated by the alternative probabilistic approach with respect to human results. Previous empirical studies on plan recognition have relied on artificially generated plans or simulated data (typically from simulation or gaming environments). Instead of testing the effectiveness of recognition algorithms in real scenarios, they mainly focus on reducing algorithms' runtime. Our work is among the first to validate plan recognition algorithm using real scenarios and by comparing with human data.

The rest of the paper is organized as follows. Section 2 reviews related work on probabilistic plan inference. Section 3 introduces probabilistic plan representation. Section 4 presents our probabilistic reasoning approach, with an illustrative example from our previous training system. We then further discuss and compare our approach with the related probabilistic models in Section 5. Section 6 conducts experimental study in group behavior prediction, including constructing plan library and test set based on the real group data and comparing prediction results of different approaches against human data. Finally, we conclude in Section 7.

2 RELATED WORK

Utility and rationality issues have been explored in AI and agent research, as means for specifying, designing and controlling rational behavior as well as descriptive means for understanding behavior. In our approach, we use utilities to represent the presumed preferences of the observed agents. State preferences are used in recognizing agents' intentions and for disambiguation. Previous research also identifies the properties of intention in practical reasoning, and states intentions as elements in agents' stable partial plans of action structuring present and future conduct. Plans thus provide context in inferring intentions, pertaining to the goals and reasons of an agent's behavior. This justifies plan inference as a means to recognize intentions and goals of agents.

In AI literature, there is a wealth of computational work on plan/intention recognition. Here we only list the most relevant work. To deal with uncertainty inherent in plan inference, Charniak and Goldman [3] build the first probabilistic model of plan recognition based on Bayesian reasoning. Huber *et al.* [4] use *PRS* as a general specification language, and construct the dynamic mapping from *PRS* to belief networks. Because of the similarity between plan recognition and natural language parsing, Panadath and Wellman [5] propose a probabilistic reasoning method based on the probabilistic state-dependent grammars (*PSDG*s). Bui *et al.* [6] propose an online probabilistic policy recognition method based on the abstract hidden Markov model (*AHMM*). More recently, Avrahami-Zilberbrand and Kaminka [7] present a hybrid approach that combines a symbolic plan recognizer with a decision-theoretic inference mechanism to capture the observer's own biases and preferences. Geib and Goldman [8] present a probabilistic plan recognition algorithm based on a plan execution model.

Though the approaches differ, most plan recognition systems infer a hypothesized plan from observation of actions. World states and in particular, the observed agents' preferences over outcomes are rarely considered in the recognition process. On the other hand, in many real-world applications, utilities of different outcomes are already known. A planning agent usually takes into account that actions may have different outcomes, and some outcomes are more desirable than the others. Therefore, when an entity makes decisions and acts on the world, it needs to balance between different possible outcomes. In this paper, we propose a decision-theoretic approach to plan recognition and explicitly model the observed agents' state preferences in the recognition process. We also conduct experimental study to validate our approach based on agents' plans and preferences in real-world scenarios.

3 PROBABILISTIC REPRESENTATION

Plan representations are used by many intelligent systems. In a traditional plan representation, an action *A* has preconditions and effects. Action *precondition* is the state that must be made true before action execution. Action *effect* (including *conditional effect*) is the state achieved after action execution. *Antecedents* and *consequents* of conditional effects are also world states. If the antecedents of a conditional effect hold before action execution, its consequents will likely hold after action execution. Actions can be either *primitive* (i.e., directly executable by agents) or *abstract*.

In a probabilistic plan representation, the likelihood of states is represented by probability values. To represent the success and failure of action execution, an action has an *execution probability* $P_{execution}$ (i.e., the likelihood of successful action execution given the preconditions are true). An action effect can be nondeterministic (i.e., *effect probability* P_{effect}, the likelihood of the occurrence of an action effect given the corresponding action is successfully executed) and/or conditional nondeterministic (i.e., *conditional probability* $P_{conditional}$, the likelihood of the occurrence of its consequent given a conditional effect and its antecedents are true). The desirability of action effects (i.e., their positive/negative significance to an agent) is represented by *utility* values. *Outcomes* are those action effects with non-zero utility values.

A plan P is an action sequence to achieve certain intended goal(s). Representation of plans is similar to that in the probabilistic plan representation, except that we use expected utilities (EU) of plans to represent the overall benefit or disadvantage of a plan.

4 PROBABILISTIC REASONING APPROACH

Our approach is based on the fundamental *MEU* ("maximum expected utility") principle underlying decision theory, which assumes that a rational agent will adopt a plan maximizing the expected utility. The computation of expected plan utility captures two important factors. One is the desirability of plan outcomes. The other is the likelihood of outcome occurrence, represented as outcome probability. The calculation of outcome probability considers three sources of uncertainty: uncertainty in action preconditions (i.e., state probabilities), uncertainty in action execution (i.e., execution probabilities), and nondeterministic and/or conditional action effects (i.e., effect probabilities). Before presenting our computational approach, we first introduce the notations we adopt.

4.1 NOTATIONS

Let E be the evidence. Let A, e, c, o and P be an action, an action effect, a consequent of a conditional effect, an outcome and a plan, respectively. The following notations are adopted in our approach.

- *precondition*(A): precondition set of action A.
- *effect*(A): effect set of action A.
- *conditional-effect*(A): conditional effect set of action A.
- *antecedent*(e): antecedent set of conditional effect e.
- *consequent*(e): consequent set of conditional effect e.
- $P_{effect}(e \mid A)$: probability of the occurrence of its effect e given action A is successfully executed.
- $P_{conditional}(c \mid \text{antecedent}(e), e)$: probability of the occurrence of its consequent c given conditional effect e and its antecedents are true.
- $P_{execution}(A \mid \text{precondition(A)})$: probability of successful execution of action A given its preconditions are true.
- $P_{action}(o|E)$: probability of action outcome o given evidence E.
- $P_{plan}(o|E)$: probability of plan outcome o given evidence E.
- *utility*(e): utility value of effect e (ranging between -100 and $+100$ in the model).
- $EU(A|E)$: expected utility of action A given evidence E.
- $EU(P|E)$: expected utility of plan P given evidence E.

4.2 COMPUTATION

The computation of expected plan utility is similar to that in decision-theoretic planning, using the utilities of outcomes and the probabilities with which different outcomes occur. In our approach, however, we use the observed evidence to incrementally update state probabilities and the probabilities of action execution, and compute an exact utility value rather than a range of utility values as in decision-theoretic planning. This is done through recursively using plan knowledge.

4.2.1 PROBABILITY OF STATES

Let E be the evidence. If state x is observed, the probability of x given E is *1.0*. Observations of actions change the probabilities of states. If action A is observed executing, the probability of each precondition of A should be *1.0*, and the probability of each effect of A is the multiplication of its execution probability and effect probability. If A has conditional effects, the probability of a consequent of a conditional effect of A is the product of its execution probability, conditional probability and the probabilities of each antecedent of the conditional effect.

- IF $x \in$ precondition(A), $P(x \mid E) = 1.0$
- IF $x \in$ effect(A), $P(x \mid E) = P_{execution}(A \mid precondition(A)) \times P_{effect}(x \mid A)$
- IF $x \in$ consequent(e) $\wedge$ $e \in$ conditional-effect(A),

  $$P(x \mid E) = P_{execution}(A \mid precondition(A)) \times P_{conditional}(x \mid antecedent(e), e) \times \prod_{e' \in antecedent(e)} P(e' \mid E)$$

If an action A is observed executed, the probability of successful execution of A given E is *1.0*. In this case, the computation above can be simplified:

- IF $x \in$ precondition(A), $P(x \mid E) = 1.0$
- IF $x \in$ effect(A), $P(x \mid E) = P_{effect}(x \mid A)$
- IF $x \in$ consequent(e) $\wedge$ $e \in$ conditional-effect(A),

  $$P(x \mid E) = P_{conditional}(x \mid antecedent(e), e) \times \prod_{e' \in antecedent(e)} P(e' \mid E)$$

Otherwise, the probability of x given E is equal to the prior probability of x.

4.2.2 PROBABILITY OF ACTION EXECUTION

If an action A is observed executed, the probability of successful execution of A given E is *1.0*, that is, $P(A|E)=1.0$. If A is observed executing, $P(A|E)$ equals to its execution probability. Otherwise, the probability of successful execution of A given E is computed by multiplying the execution probability of A and the probabilities of each action precondition.

$$P(A \mid E) = P_{execution}(A \mid precondition(A)) \times \prod_{e \in precondition(A)} P(e \mid E)$$

So the changes of state probabilities affect the probability calculation of action preconditions, and the probabilities of action execution are changed accordingly.

4.2.3 OUTCOME PROBABILITY AND EXPECTED UTILITY OF ACTIONS

The probability changes of action execution impact the calculation of outcome probabilities and expected utilities of actions. Let O_A be the outcome set of action A, and outcome $o_i \in O_A$. The probability of o_i given E is computed by multiplying the probability of A and the effect probability of o_i.

$$P_{action}(o_i \mid E) = P(A \mid E) \times P_{effect}(o_i \mid A)$$

If o_i is the consequent of conditional effect e of A, the formula above should also include the probabilities of each antecedent of the conditional effect.

$$P_{action}(o_i \mid E) = P(A \mid E) \times P_{conditional}(o_i \mid antecedent(e), e) \times \prod_{e' \in antecedent(e)} P(e' \mid E)$$

The expected utility of A given E is computed using the utilities of each action outcome in A and the probabilities with which each outcome occurs.

$$EU(A \mid E) = \sum_{o_i \in O_A} (P_{action}(o_i \mid E) \times Utility(o_i))$$

4.2.4 OUTCOME PROBABILITY OF PLANS AND EXPECTED PLAN UTILITY

Similarly, the probability changes of action execution impact the calculation of outcome probabilities and expected plan utilities. Let O_P be the outcome set of plan P, and outcome $o_j \in O_P$. Let $\{A_1, \ldots, A_k\}$ be the partially ordered action set in P leading to o_j, where o_j is an action effect of A_k. The probability of o_j given E is computed by multiplying the probabilities of each action leading to o_j and the effect probability of o_j (Note that $P(A_i|E)$ is computed according to the partial order of A_i in P).

$$P_{plan}(o_j \mid E) = (\prod_{i=1,\ldots,k} P(A_i \mid E)) \times P_{effect}(o_j \mid A_k)$$

If o_j is the consequent of conditional effect e of A_k, the formula above should also include the probabilities of each antecedent of the conditional effect.

$$P_{plan}(o_j \mid E) = (\prod_{i=1,\ldots,k} P(A_i \mid E)) \times P_{conditional}(o_j \mid antecedent(e), e) \times (\prod_{e' \in antecedent(e)} P(e' \mid E))$$

The expected utility of P given E is computed using the utilities of each plan outcome in P and the probabilities with which each outcome occurs.

$$EU(P \mid E) = \sum_{o_j \in O_p} (P_{plan}(o_j \mid E) \times Utility(o_j))$$

The intention recognition algorithm works on a possible plan set that is a subset of the plan library. Each plan in the possible plan set includes some or all of the observed actions/states. The algorithm calculates the expected utilities of each possible plan; the one with the highest expected utility is inferred as the current hypothesized plan.

4.3 ILLUSTRATION

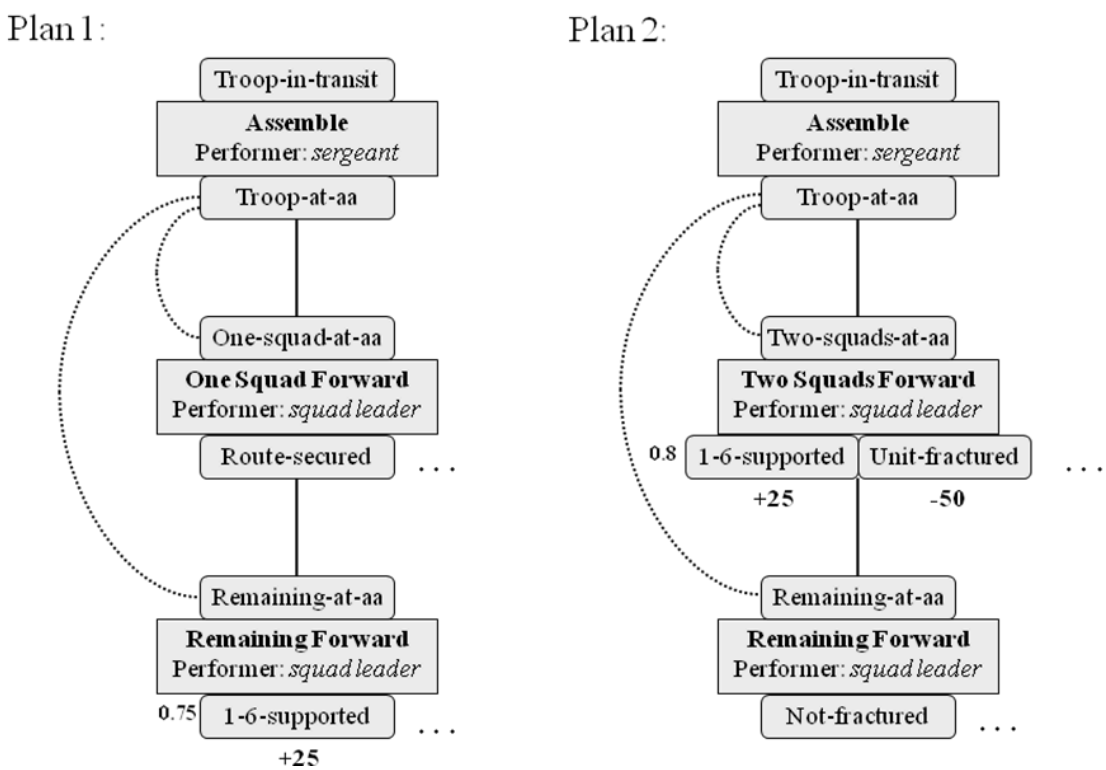

Figure 1. Two Complete Group Plans

To illustrate how our approach works, we use an example from the Mission Rehearsal Exercise (*MRE*) leadership training system. A troop led by the *lieutenant* has the mission of supporting a sister unit (i.e. unit 1-6). Three agents, the *lieutenant*, the *sergeant* and the *squad leader* act as a group in this example. The sergeant acts as an assistant of the lieutenant, and the squad leaders act as subordinates of the sergeant. *Figure 1* illustrates two complete plans of the group.

In the example, the troop's mission is to support unit 1-6. This (i.e., *unit 1-6 supported*) is a desirable goal of the group. Two plan alternatives in the plan library are available to achieve this goal, namely *Plan 1* and *Plan 2*. *Plan 1* is composed of three primitive actions, *assemble*, *one squad forward* and *remaining (squads) forward*, performed by different agents. *Remaining (squads) forward* in *Plan 1* achieves *(unit) 1-6 supported* (with effect probability *0.75*). *Plan 2* consists of primitive actions *assemble*, *two squads forward* and *remaining (squads) forward*. *Two squads forward* in *Plan 2* achieves *(unit) 1-6 supported* (with effect probability *0.8*), but also brings about undesirable outcome *unit fractured*. The corresponding effect probabilities (that are less than 1.0) and non-zero utilities are shown in the figure. The execution probability of each action is set to *0.95*.

Two actions of the group are observed executed, *assemble* and *1st-and-4th-squads-forward* (an instance of *two-squads-forward*). Based on the observations

$P(Assemble|E) = P(Two\text{-}squads\text{-}forward|E) = 1.0$

We have

$P(Troop\text{-}at\text{-}aa|E) = P(One\text{-}squad\text{-}at\text{-}aa|E) = P(Remaining\text{-}at\text{-}aa|E) = 0.95$
$P(Troop\text{-}in\text{-}transit|E)=P(Two\text{-}squads\text{-}at\text{-}aa|E) = 1.0$

Now compute the probabilities of executing *One-squad-forward* and *Remaining-forward* given the evidence. We have

$P(One\text{-}squad\text{-}forward|E) = P(Remaining\text{-}forward|E) = 0.95\times 0.95 = 0.9025$

From *Figure 1*, we know

$P_{effect}(1\text{-}6\text{-}supported|Remaining\text{-}forward) = 0.75$
$P_{effect}(1\text{-}6\text{-}supported|Two\text{-}squads\text{-}forward) = 0.8$

Now compute the probabilities of plan outcomes (Note that in *Plan 2*, outcome *unit-fractured* is deleted in the third plan step)

$P_{plan}(1\text{-}6\text{-}supported|E) = 0.9025\times 0.9025\times 0.75 = 0.61$ (*Plan 1*)
$P_{plan}(1\text{-}6\text{-}supported|E) = 0.8$ (*Plan 2*)
$P_{plan}(unit\text{-}fractured|E) = 1.0$
$P_{plan}(not\text{-}fractured|E) = 0.95$

Now compute the expected utilities of *Plan 1* and *Plan 2* using utility functions

$EU(Plan_1|E) = 0.61\times 25 = 15.25$

$EU(Plan_2|E) = 0.8\times 25+1\times (-50)+0.95\times 50 = 17.5$

The results support *Plan 2*, so *Plan 2* is recognized as the current hypothesized plan. For the more complex cases including conditional probabilities and/or abstract actions, the plan inference mechanism can be applied the same way.

5 DISCUSSIONS

Plan recognition can be characterized according to the role of the observed agent(s). The observed agent does not attempt to impact the recognition process, as if the recognizer observes the agent through a "keyhole" (i.e., *keyhole recognition*), or the observed agent deliberately performs actions to help or thwart the recognition (*intended or adversarial recognition*). The latter is typical in cooperative or competitive environments. Among the three kinds of plan recognition, keyhole recognition is the most common, without any assumption of the recognized agent's role in the recognition process. Our focus in this paper is keyhole recognition. Below we restrict our discussions to the probabilistic approaches in keyhole plan recognition.

Some probabilistic approaches have considered the influence of world states on plan recognition, when actions themselves are unobservable. For example, Bui *et al.* [6] use a variant of hidden Markov model for online policy recognition. We did not adopt a Markov model in our work for several considerations. A Markov-based approach generates relatively large state space, and assumes fixed goals. The core technologies of our application domain center on a common representation of plan knowledge, which is shared and reused among different system components. Besides, in modeling the dynamics of behavior evolvement in intelligent entities, we would like our system to give agents the flexibility of varying their interpretation of outcome desirability under different socio-cultural context.

Some work has implicitly considered an entity's utility functions. For example, Pynadath and Wellman [5] capture the likelihood that an agent will expand a plan in a particular way (i.e., the expansion probabilities of *PSDGs*). Avrahami-Zilberbrand and Kaminka [7] explicitly take the observer's preferences into consideration. Since their work is concerned with how to bias hypotheses that are more important or costly to the recognizing agent, they take the perspective of the observer (rather than the observed agents) and the utility functions are those of the observer. To address the observer's preferences, they build another decision-theoretic layer on top of the symbolic plan recognizer. In contrast, our approach takes the observed agents' state preferences and lets them participate in the recognition process.

As plan recognition involves abduction, most existing plan recognition work has employed Bayesian reasoning as a computational means to reasoning to a best explanation [3, 4, 5, 8]. This is generally realized via Bayesian (belief) networks. Although Bayesian reasoning is advantageous in accounting for how well the observed actions support a hypothesized plan and it also provides a convenient way to compute and rank different hypotheses by their probability values, the inference itself requires large numbers of prior and conditional probabilities. In many real world situations, these probabilities are hard to obtain, and there is no good answer for where the numbers come from.

Knowledge about actions, their preconditions and effects is typically available as plan representation in intelligent systems. Compared with other related work, our approach makes better use of this knowledge. Our approach also needs prior probabilities, that is, prior probabilities of states, (successful) action execution and action effects. Effect probabilities, including non-deterministic and/or conditional effects are already available in many systems with a planning component. The implications of prior state probability and probability of action execution are more intuitive, thus they are relatively easier to obtain compared to the CPTs required in Bayesian networks. Our approach incrementally uses knowledge and observations to change state probabilities and impact action/plan execution and outcome achievement.

Geib and Goldman [8] points out some limitations in current plan recognition systems. In our approach, plan library is composed of multiple partial-order plans with temporal constraints. There is no strong assumption about the observability of actions or states in our approach. Online recognition can be processed by our system in the same fashion. However, currently our system does not support interleaved goal/plan recognition (which allows agents to pursue multiple plans at a time). In addition, similar to most other recognizers, our approach treats goals as propositional and so our system does not support instantiated goal recognition. Nonetheless, we feel our approach is sufficient for the practical applications and compatible with human intuitions in recognizing intentions and goals of other entities.

Since Bayesian reasoning is the representative computational approach used by most plan recognition systems, in the next section, we empirically compare our approach with Bayesian networks in predicting group behavior using real scenarios. We report the experimental results to show the capability of our approach in modeling human intention inference in such context.

6 EXPERIMENTAL STUDY

We take three steps to evaluate the effectiveness of our approach. First, as intention recognition relies on a plan library indicating plan knowledge and recipes of plans, we construct a domain plan library using online group data. Second, based on the randomly generated evidence set and the plan library, human raters help build the test set by providing predictions associated with each line of evidence. At last, to validate our approach, we compare our model predictions with the prediction results by Bayesian reasoning against human predictions.

6.1 Construction of Plan Library

We conduct our experiment in security informatics domain and choose *Al-Qaeda* as a representative radical group for our study. Group plans can be written out manually by domain experts. However, due to the workload of hand-made plans, inconsistency between different experts and complexity of group behavior, this method is impractical and error-prone in practice. As huge volume of reports about this group and its historical events are available online, we employ computational methods to automatically generate group attack plans from relevant open source textual data [2].

The textual data we use are the news about *Al-Qaeda* reported from 2000 to 2009 in *Times Online* and *USATODAY*, with totally 10419 Web pages. Group actions are acquired by extracting verb-object pairs in each sentence where the subject is the name of the group. We design a number of linguistic patterns and use syntax parsing to extract knowledge of action preconditions and action effects for the automatic construction of domain theory [2]. The extracted group actions are then refined by unifying syntactic forms, combining semantically similar pairs and eliminating static verbs and low frequency ones, all referencing the WordNet. The refinements of action preconditions and effects are performed similarly. We finally collect 503 group actions, 110 action preconditions and 60 action effects with quality [2].

One major difficulty of domain knowledge extraction is that some of the commonsense knowledge is seldom mentioned explicitly in online news. For example, action *get visa* has the effect *have visa*, but this piece of knowledge is hard to obtain online. We compensate for missing preconditions and effects associated with group actions by adding commonsense knowledge of the verbs in action description. With the complete domain theory, we then employ planning algorithm to automatically generate attack plans of the group [2].

Among the official investigation reports, 13 real attacks perpetrated by *Al-Qaeda* have relatively complete descriptions. Based on our automatically generated plans, intelligence analyst helped choose 13 plans that match the reported real attacks. These plans form the plan library for our experimental study. Another consideration is that although using large numbers of plans is computationally feasible by our approach, we would prefer a relatively small and realistic plan library so that it is tractable by human raters in the experiment.

Figure 2 shows the example of a real attack plan claimed by *Al-Qaeda*. The corresponding action knowledge, action execution probabilities, effect probabilities and utilities are also given below. Outcome utilities in these plans are the normalized values calculated based on the *GTD* (*Global Terrorism Database*) data of the reported real or estimated damage (in the cases of success or failed attempt) of the actual attacks by this group in history (The assumption here is that causing loss or damage is desirable to this group). The average length of the plans in the plan library is *9.8* (including start and end actions).

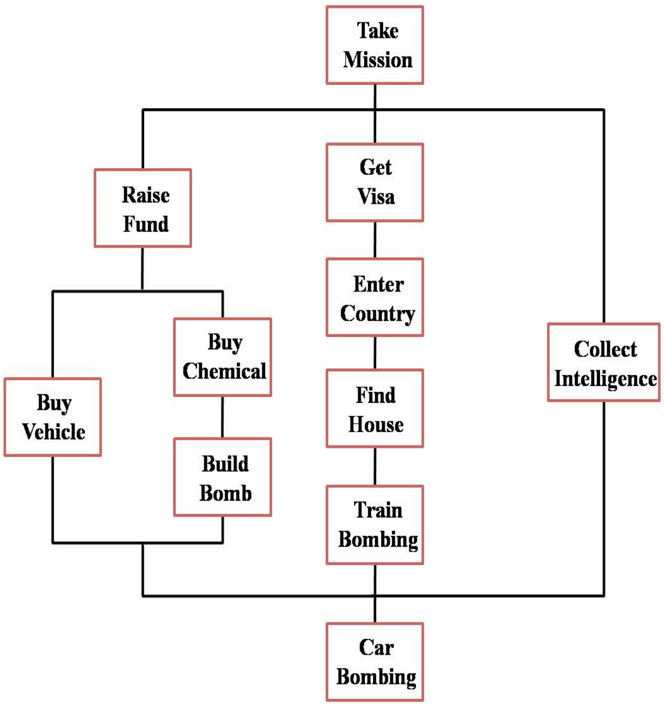

Action: take mission (1.0)
  Precondition: have authority, have operator
  Effect: have mission (1.0)
Action: raise fund (0.75)
  Precondition: have mission, have operator
  Effect: have money (0.5)
Action: buy vehicle (0.95)
  Precondition: have operator, have money
  Effect: have vehicle (0.95)
Action: buy chemical (0.75)
  Precondition: have operator, have money
  Effect: have chemical (0.5)
Action: build bomb (0.75)
  Precondition: have operator, have chemical
  Effect: have bomb (0.75)
Action: get visa (0.95)
  Precondition: have mission, have operator
  Effect: have visa (0.5)
Action: enter country (0.95)
  Precondition: have operator, have visa
  Effect: in country (0.95)
Action: find house (0.95)
  Precondition: have operator, in country
  Effect: have house (0.75)
Action: train bombing (0.95)
  Precondition: have operator, have house
  Effect: master bombing (0.95)
Action: collect intelligence (0.75)
  Precondition: have mission, have operator
  Effect: have intelligence (0.5)
Action: car bombing (0.5)
  Precondition: have vehicle, have bomb,
  master bombing, have intelligence
  Effect: have damage (0.1)
  Utility: 11

Figure 2. A Group Attack Plan in the Plan Library

6.2 The Test Set

We randomly generate a set of evidence using the combination of actions and initial world states in the plan library. We classify these actions and states into eight groups, in which similar action/state pairs or mutually exclusive actions are grouped together. For example, {Action: *raise fund*, State: *have money*} and {Action: *buy vehicle*, State: *have vehicle*} are similar action/state pairs, belonging to *Group 1* and *Group 4* respectively; {Action: *buy bomb*, Action: *build bomb*} and {Action: *take plane*, Action: *take train*} are mutually exclusive action sets, belonging to *Group 5* and *Group 7*, respectively. *Group 8* contains those last actions in each plan, such as *plane attack*, *plane bombing*, *train bombing*, *car bombing*, *suicide bombing*, *bomb attack*, and *shoot attack*. In order to obtain meaningful evidence set, the random generation process only selects actions and/or states from different groups.

However, we did not use *Group 8* in generating the evidence set. These last actions in plans are much more indicative (which associate with the goals of each plan), therefore our approach always yields very good results given them. We also delete the generated evidence with conflict actions or action/state pairs, e.g. *train hijack* and *take train*. We finally collect *95* lines of evidence. Each line contains either two observations (*49%* of the evidence set) or three observations (*51%* of the evidence set).

Four human raters experienced in security informatics participate in the experiment. Based on the plan library we construct, each rater examined the evidence set line by line and predicted the most likely plans associated with each line of evidence. The test set is composed of their predictions based on the evidence, with inter-rater agreement (*Kappa*) *0.764*. The prior state probabilities, action execution probabilities and effect probabilities used by our approach (less than *100* items in total) were assigned by intelligence analyst. The intelligence analyst also assigned prior and conditional probabilities for Bayesian reasoning. Mapping plans to Bayesian networks is based on the generic method provided in [4].

6.3 RESULTS

Table 1 shows the experimental results using our approach and Bayesian reasoning. We measure the agreement of our approach and each rater using the *Kappa statistic*. The *Kappa* coefficient is the de facto standard to evaluate the agreement between raters, which factors out expected agreement due to chance. The average agreement between our approach and human raters is *0.664* (for *two observations*) and *0.773* (for *three observations*), which significantly outperform the average agreement between Bayesian reasoning and the raters. As $0.6<k<0.8$ indicates substantial agreement, the empirical results show good consistency between the predictions generated by our approach and those of human raters.

The results also show that compared to Bayesian reasoning, the performance of our approach improves rapidly with the increase of the number of evidence. As our approach makes use of action knowledge in the inference process, actions and states are closely connected and the change in one action or state will quickly propagated to the other interrelated actions and/or states. Thus our approach is more sensitive to the number of observations. For four observations, there is unanimous agreement between human raters, and our approach shows excellent agreement with the raters (i.e., *convergence point*). As our algorithm only considers a subset of plans in the plan library that are consistent with current observations (that is, each possible plan being considered includes at least one observed action or state), to be fair in comparison, we apply Bayesian reasoning in the same way (This has increased the average agreement of Bayesian reasoning and human raters from average *0.25* to the values in *Table 1*).

Table 1 *Kappa* Agreements between Algorithms and Human Raters

In addition, we compare the answers of human raters and *1-best* and *2-best* results of the algorithms. We find that high percentage of the *2-best* results generated by our algorithm falls into the raters' answers (see *Table 2*). Bayesian reasoning also improves considerably in the *2-best* case. The average percentages of the total match of our approach are *75%* for *1-best* and *90.26%* for *2-best*.

Table 2 Comparison of Algorithms' *1-Best* and *2-Best* Results and Rater Answers

Rater	Two Observations			Three Observations			Two Observations			Three Observations		
	P(A)	***P(E)***	***K***	***P(A)***	***P(E)***	***K***	***P(A)***	***P(E)***	***K***	***P(A)***	***P(E)***	***K***
1	0.826	0.108	0.805	0.878	0.106	0.861	0.436	0.078	0.388	0.469	0.106	0.406
2	0.696	0.090	0.666	0.837	0.105	0.818	0.435	0.086	0.382	0.469	0.107	0.405
3	0.609	0.096	0.567	0.776	0.097	0.752	0.435	0.098	0.374	0.490	0.102	0.432
4	0.652	0.088	0.618	0.694	0.101	0.660	0.370	0.089	0.308	0.388	0.092	0.326
AVG			**0.664**			**0.773**			**0.363**			**0.392**

7 CONCLUSION

This paper presents a decision-theoretic approach to plan inference based on the principle of maximizing expected plan utility. Our approach considers both actions and states in the recognition process, and explicitly takes the observed agent's preferences into consideration. Online plan recognition is realized by incrementally using plan knowledge and observations to change state probabilities and impact action/plan execution and outcome achievement.

Rater	Plan Inference				Bayesian Reasoning			
	1-Best		***2-Best***		***1-Best***		***2-Best***	
	#Match	*#Error*	*#Match*	*#Error*	*#Match*	*#Error*	*#Match*	*#Error*
1	81	14	91	4	43	52	58	37
2	73	22	85	10	43	52	56	39
3	66	29	84	11	44	51	64	31
4	65	30	83	12	36	59	57	38
Percent-age	**75%**	25%	**90.26%**	9.74%	**43.68%**	56.32%	**61.84%**	38.16%

Group behavior prediction as an emergent research and application field is gaining increasing attention in recent years. It provides an ideal testbed for practicing and evaluating plan inference approaches. Based on the realistic online group data, we construct plan library and conduct experiment to evaluate our approach in group behavior prediction. We empirically compare our work with the alternative probabilistic approach. The experimental results show that our approach is compatible with human intuitions in intention/goal recognition and effective in practice.

In addition to group behavior prediction, we believe our approach is applicable to a wide range of fields. Our future research will exploit this work in several ways, including extracting and analyzing group behavior patterns from online social media, modeling organization behavior in artificial society [9], and extending intention recognition of entities for social inference and social computing [10].

ACKNOWLEDGEMENTS

This work is supported in part by the NNSFC under Grant Nos. 60875028 and 60921061, and the MOST of China under Grant No. 2006AA010106.

REFERENCES

[1] V. Martinez, G. I. Simari, A. Sliva and V. S. Subrahmanian. CONVEX: Similarity-Based Algorithms for Forecasting Group Behavior. IEEE Intelligent Systems, 23(4):51-57, July/August, 2008.

[2] X. Li, W. Mao, D. Zeng and F. Wang. Automatic Construction of Domain Theory for Attack Planning. Proceedings of the 2010 IEEE International Conference on Intelligence and Security Informatics, pp.65-70. IEEE Press, 2010.

[3] E. Charniak and R. Goldman. A Bayesian Model of Plan Recognition. Artificial Intelligence, 64(1):53-79, 1993.

[4] M. J. Huber, E. H. Durfee and M. P. Wellman. The Automated Mapping of Plans for Plan Recognition. Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence, 1994.

[5] D. V. Pynadath and M. P. Wellman. Probabilistic State-Dependent Grammars for Plan Recognition. Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence, 2000.

[6] H. H. Bui, S. Venkatesh and G. West. Policy Recognition in the Abstract Hidden Markov Model. Journal of Artificial Intelligence Research, 17:451-499, 2002.

[7] D. Avrahami-Zilberbrand and G. A. Kaminka. Incorporating Observer Biases in Keyhole Plan Recognition. Proceedings of the Twenty-First National Conference on Artificial Intelligence, pp.944-949. AAAI Press, 2007.

[8] C. W. Geib and R. P. Goldman. A Probabilistic Plan Recognition Algorithm based on Plan Tree Grammars. Artificial Intelligence, 173:1101-1132. Elsevier, 2009.

[9] F. Wang. The Emergence of Intelligent Enterprises: From CPS to CPSS. IEEE Intelligent Systems, 25(4):85-88, July/August, 2010.

[10] F. Wang, K. M. Carley, D. Zeng and W. Mao. Social Computing: From Social Informatics to Social Intelligence. IEEE Intelligent Systems, 22(2):79-83, March/April, 2007.

BIOGRAPHY

Wenji Mao received her Ph.D. degree in computer science from the University of Southern California in 2006. She is an associate professor at the Institute of Automation, Chinese Academy of Sciences. Prof. Mao is a member of ACM and AAAI. Her research interests include artificial intelligence, intelligent agents and social modeling. Contact her at wenji.mao@ia.ac.cn.

Jonathan Gratch received his Ph.D. degree in computer science from the University of Illinois at Urbana-Champaign in 1995. He is a research associate professor at the Computer Science Department, University of Southern California and the associate director of virtual human research at USC Institute for Creative Technologies. His research interests include artificial intelligence, cognitive science, emotion modeling, planning and multi-agent systems. Contact him at gratch@ict.usc.edu.

Xiaochen Li is a Ph.D. candidate at the Institute of Automation, Chinese Academy of Sciences. His research interests mainly focus on social computing and agent-based modeling. Contact him at xiaochen.li@ia.ac.cn.