Notes for Meeting 23
Plan Understanding

Understanding Others' Behavior

Although most AI research on plans and action has focused on generation, humans are also able to understand other's actions.

This capacity, sometimes called plan understanding, plays a key role in our ability to interact with others.

Plan understanding lets us construct models of other agents' mental states based on their observed activities.

Intelligent agents that lack this ability are doomed to treat other agents as objects without beliefs, goals, or intentions.

Applications of Plan Understanding

The task of plan understanding arises in many different situations:

- socially-constrained physical activities like driving and walking
- competitive activities like sports, board games, and computer games
- inferring enemry intentions in support of military planning
- coordinating complex joint activities with other agents
- appreciating comedies and dramas on television
- carrying out an extended dialogue with another agent

The ability to understand other agents' plans is a distinguishing feature of human intelligence.

The Task of Plan Understanding

We can define the task of plan understanding semi-formally as:

- Given: Background knowledge about actions and/or plans that can occur in some domain.
- Given: An observed sequence of states traversed / actions taken by another agent.
- Find: One or more plans that account for the observed behavior in terms of background knowledge.

This task is substantially more challenging that activity recognition, which involves assigning observed behavior to some activity class.

Forms of Knowledge for Plan Understanding

Methods for plan understanding must encode background knowledge in some form, such as:

- the conditions and effects of domain actions
- hierarchical structures like HTN methods
- specific plans that have occurred in the past

In general, richer forms of background knowledge make the task of plan understanding more tractable.

Three Paradigms for Plan Understanding

AI researchers have explored three main approaches to interpreting other agents' plans:

- Deductive approaches (e.g., Kautz & Allen, 1986), which aim to prove the agent was pursuing one of a set of plans.

  - These support rich representations but require enough knowledge to eliminate most alternatives.

- Probabilistic approaches, which attempt to find the most probable plan given the observations.

  - These typically rely on variants of hidden Markov models and assume limited representations.

- Abductive approaches, which aim to find one or more plausible acounts of the agent's behavior.

  - These support rich representations but do not require probabilities and do not need to disprove alternatives.

Both the deductive and abductive frameworks can benefit from structural knowledge, such as HTN methods.

Understanding Multi-Agent Interactions

Some settings require one to understand observed behavioral interactions that occur among multiple agents.

These require the ability to represent and reason about more embedded structures than single-agent behavior, such as:

- (goal Joe (loves Mary Joe))
- (belief Joe (goal Mary (loves Sam Mary))
- (goal Mary (loves Joe Mary))

An agent observing Joe, Mary, and Sam must infer not only their beliefs and goals, but their beliefs about each others' belief and goals.

Misunderstandings are the basis for many comedies and tragedies, which we appreciate because we can interpret them.

Understanding and Learning

Plan understanding also provides a source of material to drive learning.

- Given: Background knowledge about actions and/or plans that can occur in some domain.
- Given: An explanaiton for an observed sequence of states traversed / actions taken by another agent.
- Find: New hierarchical skills or methods that let one interpret similar situations or generate similar behavior.

Such analytical approaches to learning acquire knowledge far more rapidly than empirical / statistical methods.

Assignments for Meeting 24
Cognition, Affect, and Emotion

Read the article:

* Marsella, S., Gratch, J., & Petta, P. (in press). Computational models of emotion. In Scherer, K. R., Bnziger, T., & Roesch, E. (Eds.), A blueprint for an affectively competent agent: Cross-fertilization between emotion psychology, affective neuroscience, and affective computing. Oxford: Oxford University Press. [required]

* Simon, H. A. (1967). Motivational and emotional controls of cognition. Psychological Review, 74, 29-39. [optional]

* Marinier, R., & Laird, J. (2004). Towards a comprehensive computational model of emotions and feelings. Proceedings of the Sixth International Conference on Cognitive Modeling. Pittsburgh, PA. [optional]