Modeling Reciprocity in Human Bilateral Negotiation

Ya'akov Gal
Avi Pfeffer

Motivation

Goal: learn how psychological factors affect human negotiation in environments that include

- hierarchy between goals, tasks and resources
- strategic complexity
- social relationships

Method:

- formally represent factors within a computational framework
- learn a hierarchical model from observing people
- evaluate model's ability to
  - predict behavior
  - generalize across different types of people and environments

Colored Trails (Grosz and Kraus '04)

A framework for investigating decision-making between computers and/or people. Interesting for people to play; possible for computers to play

- provides analogy to task settings; abstracts away from domain specific knowledge.
- flexible and extendible. Can vary (among others)
  - complexity of interaction
  - uncertainty of observations
  - environment type (competitive/collaborative)
  - negotiation protocol

Multi-attribute Negotiation Scenario

Unknown, finite number of one-shot CT rounds

- full visibility of board and chips
- dependency relationships vary
- proposer player makes a bid
- responder accepts/rejects bid
- agreement is enforced
- individual performance depends on reaching goal, chips and path taken
- roles alternate at each round

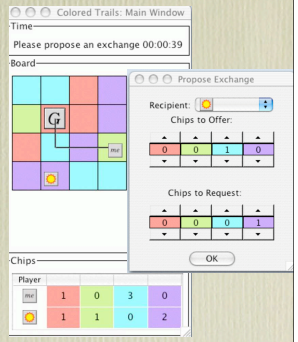

MURI relevance

Features of repeated interaction across cultures

- Belief revision (e.g., ping-pong diplomacy); must account for inconsistent/incorrect beliefs.
- Dynamic reasoning
  - prospective: what is the ramification (cost) of a potential action in an uncertain future ?
  - retrospective: what is the extent to which we reward/punish others for their actions ?
- need to consider other's intentions, their beliefs about each other's intentions, etc...

Modeling Intention

Agent *j*'s *merit* - the (normalized) difference between the "fair" outcome for *i* and the outcome *i* believes *j* will choose.

1>j's *merit* > *0* : *(i* believes that) *j* is "kind"

-1<j's *merit* < *0* : *(i* believes that) j is "nasty"

(*i*'s beliefs about) Retrospective merit for *j*: payoff for *j* * *j*'s merit

Fairness Equilibrium (Rabin '99): agents positively reciprocate kindness; negatively reciprocate nastiness.

example (prisoner's dilemma) : both players believe the other intends to be nice/nasty, and both players cooperate/ defect.

We define a fair exchange as the Nash Bargaining solution concept.

Modeling Repeated negotiation

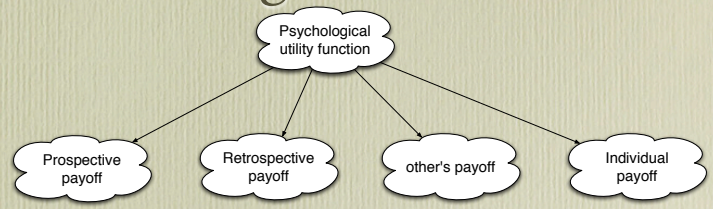

Intentions are updated by agents at each round

Assume a finite set of types; each type's utility function weighs social factors differently.

Learn a mixture model

- probability distribution over types
- for each type, learn weights for each social factors

Generative Model

Each instance (round) consists of a set of CT rounds. Each round *i* including an offer c_i and reply r_i. Each action is generated by (unobserved) type t_1,t_2.

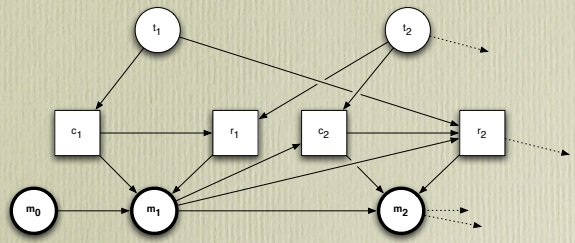

$$P(\{c_1,\ldots,r_1\},\ldots,\{c_n,\ldots,r_n\} \mid t_1,t_2) =$$
$$P(c_1 \mid t_1,\mathbf{m}_0)P(r_1 \mid c_1,t_2,\mathbf{m}_0)P(c_2 \mid t_2,\mathbf{m}_1)\ldots P(r_n \mid c_{n-1},t_2,\mathbf{m}_{n-1})$$

Computing Future ramification

Sample from a set of future games.

For each potential action, compute the expected reward from future games, given current parameter models.

Example - for responder player i, future ramification at horizon k at game g_n, merit scalars m_i, m_j is computed by

$$ED_K(r_n^i \mid g_n, p_n^j, m_i, m_j) = S_i(p_n^j, r_n^i, g_n) + \sum_{g_{n+1}} \sum_{p_{n+1}^i} P(p_{n+1}^i \mid g_{n+1}, m_j) EP_{K-1}(p_{n+1}^i \mid g_{n+1}, m_i', m_j)$$

Preliminary Results

Collected 70 instances; each includes a set of CT interactions.

Correlations suggest that reciprocal interactions are present.

- (positive) benefit to responde/response r in two consecutive rounds, response in.
- (negative) response and responder's benefit in two consecutive rounds.

Fit-to-data experiments

- predictive model that learned reciprocal behavior outperformed models that learned one shot behavior,

Future Work/Obtaining CT

Compare "social" agent performance to a variety of other agents.

- Game theoretic agents
- Other people

Evaluating mode by

- playing new people in new types of situations.
- cultural distinctions

CT release (version 3), available for download

- www.eecs.harvard.edu/ai/ct3
- email list - release announcements, bug fixes, etc...

Modular code can be tailored by modifying a Java configuration class

- Examples: single vs. repeated interaction, prescribed or