Using Gaussian Processes to Optimise Concession in Complex Negotiations against Unknown Opponents

Colin R. Williams, Valentin Robu, Enrico H. Gerding and Nicholas R. Jennings

Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science, University of Southampton

Introduction

- **Features of Complex Negotiation**
  - Alternating offers protocol
  - Multiple negotiation issues
  - Discrete and continuous issues
  - Real-time constraints
    - Discounting factor
    - Deadline
- **Unknown Opponents**
  - Unknown utility function
  - Unknown behaviour
  - Single negotiation encounter

Predicting the Opponent's Concession

1. Observe offers made by opponent
2. Take best offer seen in time window
3. Perform Gaussian process regression (repeated at the end of each time window)

Setting Concession Rate

1. Find best time, t* to reach agreement
   - By maximising the expected discounted utility of opponent's offers
2. Find best utility, u* to propose offers at
   - By maximising the expected discounted utility of our offer at time t*
3. Concede towards this point [t*, u*]
4. Repeat until agreement or deadline is reached

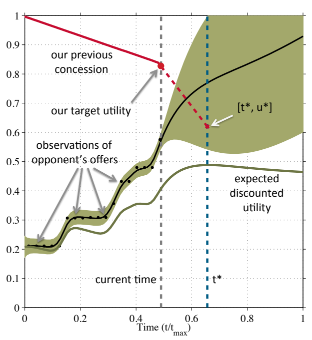

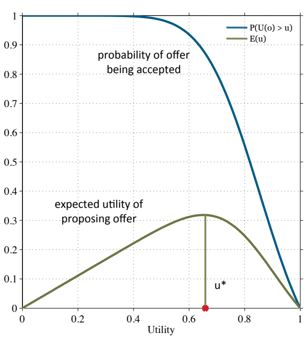

Spiteful Behaviour

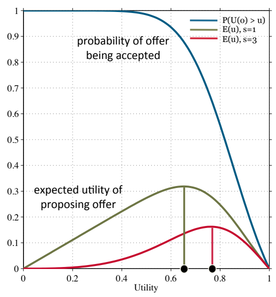

- Useful in competition environment: Automated Negotiating Agent Competition (ANAC).
- May wish to 'win' negotiation by reaching better agreements than opponents.
- Spiteful behaviour aims to avoid low utility agreements and therefore concedes less.

Evaluation

- Evaluated against ANAC 2010 negotiation agents, averaged over a variety of domains
- Tested with two risk profiles (s=1, s=4)
- Highest tournament score: 0.492 (s=1)
- Highest self-play score: 0.722 (s=1)

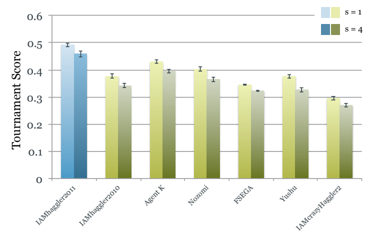

Empirical Game Theoretic Analysis

- Compare range of tournaments
- Find incentives for single agent to change strategy
- Search for equilibria

Each node represents a tournament of six agents using one of three strategies, from the discounted version of the largest domain used in ANAC 2010 (the travel domain).

Each edge represents the best single-agent deviation from one mixture to another.

Two equilibria (all agents use strategy IAMhaggler2011 or all agents use strategy IAMhaggler2010).

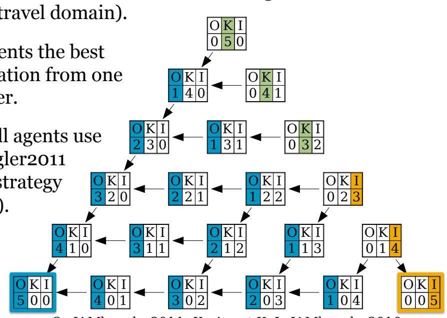

O=IAMhaggler2011, K=Agent K, I=IAMhaggler2010

Colin R. Williams, Valentin Robu, Enrico H. Gerding and Nicholas R. Jennings (2011) Using Gaussian Processes to Optimise Concession in Complex Negotiations against Unknown Opponents. *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, AAAI Press.

UNIVERSITY OF Southampton
School of Electronics and Computer Science