

Automated negotiation

Sarit Kraus

Department of Computer Science

Bar-Ilan University

University of Maryland

sarit@cs.biu.ac.il

<http://www.cs.biu.ac.il/~sarit/>

Automated Negotiation

Part 2: Automated agents that interact proficiently with people

Main Points

Agents negotiating with people is important

General opponent* modeling:

machine learning




human behavior model

Computers interacting with people

Computer
has the
control

Human has
the control





Culture sensitive agents

The development of standardized
Buyer/Seller agents in culture and

Buyer/Seller
agents
well as
culture

PURB
agent




Semi-autonomous cars



Medical applications

**Gertner Institute for
Epidemiology and Health
Policy Research**



Automated care-taker




I will be too tired in the afternoon!!!

The physiotherapist has no other available appointments this week. How about resting before the appointment?


Try to reschedule and fail

Security applications



People often follow suboptimal decision strategies

- ▶ Irrationalities attributed to
 - sensitivity to context
 - lack of knowledge of own preferences
 - the effects of complexity
 - the interplay between emotion and cognition
 - the problem of self control
 - bounded rationality in the bullet



Agents that play repeatedly
with the same person

AutONA [BY03]

- ▶ Buyers and sellers
- ▶ Using data from *previous* experiments
- ▶ Belief function to model opponent
- ▶ Implemented several tactics and heuristics
 - including, concession mechanism

A. Bye, M. Yearworth, K.-Y. Chen, and C. Bartolini. AutONA: A system for automated multiple 1-1 negotiation. In *CEC*, pages 59–67, 2003

General opponent* modeling

same person only once

Challenges of human opponent* modeling

- ▶ Small number of examples
 - difficult to collect data on people
- ▶ Noisy data
 - people are inconsistent (the same person may act differently)
 - people are diverse

Guessing Heuristic

- ▶ Multi-issue, multi-attribute, with incomplete information
- ▶ Domain independent
- ▶ Implemented several tactics and heuristics
 - including, concession mechanism

C. M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *JAAMAS*, 15(2):221–252, 2007

Single interactions:

Cliff-Edge environments (CE)

- ▶ First-price Sealed-bid Auctions
 - Possible values $[0 \dots N]$
 - Equilibrium strategy: N
- ▶ Dynamic pricing
- ▶ Ultimatum game (from proposer's point of view)

Agent plays once with each person.




Cliff-Edge

DVRL: **Reinforcement learning (RL) +
Virtual Learning (V) +
Deviation (D)**


- Using data from *previous* interactions
- Implemented several tactics and heuristics
 - qualitative in nature
- *Non-deterministic* behavior, via means of randomization

R. Katz and S. Kraus. Efficient agents for cliff edge environments with a large set of decision options.
In *AAMAS*, pages 697–704, 2006


Deviated VRL




Virtual RL




RL




Gender-Sensitive Negotiator (GSN)




Gender Sensitive



Full separator




Generic



QOAgent [LIN08]

- ▶ Multi-issue, multi-attribute, with incomplete information
- ▶ Domain independent
- ▶ Implemented several tactics and heuristics
 - qualitative in nature
- ▶ *Non-deterministic* behavior, also via means of randomization

R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7):823–851, 2008



employer and job
candidate

Objective: reach an
agreement over hiring
terms after successful
interview

QOvalue components (AIJ08)

► Luce numbers

- Attached to each offer
- Equals to the relative utility value of an offer:

$$lu(\text{offer}) = \frac{u(\text{offer})}{\sum_{\vec{o} \in O} u(\vec{o})}$$

► Offer ranking

- We have a finite set of agreements
- All agreements are ranked based on their utility value in the range (0,1]

QOvalue (motivation)

- ▶ Taking into account the opponent's point-of-view when making an offer
 - Generate **good** offers for the agent, which is estimated to be *acceptable* by the opponent
 - Estimating acceptance using Luce numbers

$$QOValue(\vec{o}) = \min\{\alpha, \beta\}$$

$$\alpha = rank_{agent}(\vec{o}) \cdot lu_{agent}(\vec{o})$$

$$\beta = [lu_{opponent}(\vec{o}) + lu_{agent}(\vec{o})] \cdot rank_{opponent}(\vec{o})$$

QOAgent (AIJ08)

- ▶ Make an offer with the highest QOvalue
- ▶ Accept an offer with a probability that depends on the ranking number of the offer

QOAgent [LIN08]

Played at least as well as people

- ▶ Multi-issue, multi-attribute, with incomplete information
- ▶ Domain independent
- ▶ Implemented several tactics and heuristics
 - qualitative in nature
- ▶ No deterministic behavior

Yes, if you have data

Is it possible to improve the QOAgent?

R. Lin, S. Kraus, J. Wilkenfeld, and S. Edelkamp. Playing a multi-issue game with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(3):83–851, 2008


KBAgent

- ▶ Multi-issue, multi-attribute, with incomplete information
- ▶ Domain independent
- ▶ Implemented several tactics and heuristics
 - qualitative in nature
- ▶ *Non-deterministic* behavior, also via means of randomization
- ▶ Using data from **previous interactions**

Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 2009

General opponent modeling

- Challenge: sparse data of past negotiation sessions of people negotiation
- Technique: *Kernel Density Estimation*



General opponent modeling

- ▶ Estimate likelihood of other party:
 - accept an offer
 - make an offer
 - its expected average utility
- ▶ The estimation is done separately for each possible agent type:
 - The type of a negotiator is determined using a simple Bayes' classifier
- ▶ **Use estimation for decision making**

KBAgent's decision: making offers

- ▶ Compute QOvalue for possible offers and propose the offer with the highest QOvalue (similar to the *QOAgent*)
- ▶ Create a concession rate policy using the opponent model for additional offers (if needed)

KBAgent - Concession Policy

- ▶ Create the list of offers:
 - with utility value above the SQ
 - discard offers that reduce both the other party's utility and the *QOvalue*.

KBAgent Concession policy

Offer Index	$QOValue(\vec{o}_t)$	$u_{opp}(\vec{o}_t)$	Order list
	10	435	
	12	400	
	12.5	300	
	9	410	
	10.3	350	
	9.87	470	
	9.8	490	
	13.45	350	
	8.8	500	
	11.22	430	

*6 possible
negotiation
rounds

*Average
expected 480

KBAgent Concession policy

Offer Index	$QOValue(\vec{o}_t)$	$u_{opp}(\vec{o}_t)$	Order list
0	13.45	350	
1	12.5	300	discard
2	12	400	
3	11.22	430	
4	10.3	350	discard
5	10	435	
6	9.87	470	
7	9.8	490	
8	9	410	discard
9	8.8	500	

*6 possible
negotiation
rounds

*Average
expected 480

KBAgent - Concession Policy

- ▶ Use a concession rate to propose offers from the list
 - decreasing *QOvalue* for the agent
 - increasing utility value for the other party
 - The concession rate is determined such that after passing 80% of the turns, the *KBAgent* will propose the offer which is as closest to the *Expected Opponent Average*

KBAgent Concession policy

Offer Index	$QOValue(\vec{o}_t)$	$u_{opp}(\vec{o}_t)$	Order list
0	13.45	350	0
1	12.5	300	-
2	12	400	1
3	11.22	430	2
4	10.3	350	-
5	10	435	3
6	9.87	470	4
7	9.8	490	5
8	9	410	-
9	8.8	500	6

*6 possible
negotiation
rounds

*Average
expected 480

Computing expected utility for acceptance threshold

$$E(t, \alpha_t) = \sum_{u_{KB}(\vec{o}, t) \geq \alpha_t} P(\vec{o}, t) u_{KB}(\vec{o}, t) +$$
$$\left(\sum_{u_{KB}(\vec{o}, t) < \alpha_t} P(\vec{o}, t) \right) \{ Q(o_{KB}(t+1)) \cdot u_{KB}(o_{KB}(t+1), t+1) +$$
$$[1 - Q(o_{KB}(t+1))] E(t+1, \alpha_{t+1}) \}$$

$O_{kb}(t)$ agreement to be offered by agent

KBAgent as the job candidate

- ▶ Best result: 20,000, Project manager, With leased car; 20% pension funds, fast promotion, 8 hours

Human



12,000
Programmer
Without leased car
Pension: 10%
Fast promotion
10 hours

KBAgent



20,000
Team Manager
With leased car
Pension: 20%
Slow promotion
9 hours

20,000
Project manager
Without leased car
Pension: 20%
Slow promotion
9 hours


KBAgent as the job candidate

- Best agreement: 20,000, Project manager, With leased car; 20% pension funds, fast promotion, 8 hours

12,000
Programmer
Without leased car
Pension: 10%
Fast promotion
10 hours

20,000
Team Manager
With leased car
Pension: 20%
Slow promotion
9 hours

KBAgent



Round 7



20,000
Programmer
With leased car
Pension: 10%
Slow promotion
9 hours

Experiments

**Learned from 20 games
of human-human**

- ▶ 172 grad and undergrad students in Computer Science
- ▶ People were told they may be playing a computer agent or a person.
- ▶ Scenarios:
 - Employer-Employee
 - Tobacco Convention: England vs. Zimbabwe

Results:

Comparing *KBAgent* to others

Player	Type	Average Utility Value (std)
<i>KBAgent vs people</i>	Employer	468.9 (37.0)
<i>QOAgent vs peoples</i>		417.4 (135.9)
<i>People vs. People</i>		408.9 (106.7)
<i>People vs. QOAgent</i>		431.8 (80.8)
<i>People vs. KBAgent</i>		380. 4 (48.5)
<i>KBAgent</i>	Job Candidate	482.7 (57.5)
<i>QOAgent</i>		397.8 (86.0)
<i>People vs. People</i>		310.3 (143.6)
<i>People vs. QOAgent</i>		320.5 (112.7)
<i>People vs. KBAgent</i>		370.5 (58.9)

Main results

- ▶ In comparison
 - The KP values than
 - QO
 - M
 - higher
- ▶ T...er utility
- ▶ Resu... negotiation
done by th...

General
opponent
modeling
improves agent
negotiations

Task

The development of standardized agent to be used in the collection of data for structure and negotiation

Simple
Computer
System



Motivation

- ▶ Technology has revolutionized communication
 - Cheap and reliable
 - Transcends geographic boundaries
- ▶ People's cultural background significantly affects the way they communicate
- ▶ For computer agents to negotiate well across cultures they need to be highly adaptive to behavioral traits that are culture-specific

KBAgent [os09]

- ▶ Multi-issue, multi-agent, incomplete information
- ▶ Domain independent
- ▶ Implementation characteristics
 - qualitative
- ▶ *Non-deterministic* randomization
- ▶ Using data from **previous interactions**

No previous data

Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 2009




QOAgent [LIN08]

- ▶ Multi-issue, multi-attribute, with incomplete information
- ▶ Domain independent
- ▶ Implemented several tactics and heuristics
 - qualitative in nature
- ▶ *Non-deterministic* behavior, also via means of randomization


R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7):823–851, 2008

GENIUS interface



R. Lin, S. Kraus, D. Tykhonov, K. Hindriks and C. M. Jonker.
Supporting the Design of General Automated Negotiators.
In ACAN 2009.

Example scenario



Employer and job
candidate

reach an
agreement over hiring

with



Cliff-Edge [KA06]

- Repeated ultimatum game
- Virtual learning and reinforcement learning
- Gender-se


Too simple
scenario;
well studied

R. Katz and S. Kraus. Efficient
for cliff edge environments with a
set of decision options. In *AAMAS*,
pages 697–704, 2006




Color Trails (CT)

- ▶ An infrastructure for agent design, implementation and evaluation for open environments
- ▶ Designed with Barbara Grosz (AAMAS 2004)
- ▶ Implemented by Harvard team and BIU team



An Experimental Test-Bed

- ▶ Interesting for people to play:
 - analogous to task settings;
 - vivid representation of strategy space (not just a list of outcomes).
- ▶ Possible for computers to play.
- ▶ Can vary in complexity
 - repeated vs. one-shot setting;
 - availability of information;
 - communication protocol.



Board Display

		G		
me				O

Player Chips Display

me	5	0	10	0	9
O	1	0	13	10	0

Scoring and payment

- ▶ 100 point bonus for getting to goal
- ▶ 10 point bonus for each chip left at end of game
- ▶ 15 point penalty for each square in the shortest path from end-position to goal
- ▶ Performance does not depend on outcome for other player

Colored Trails: Motivation

- ▶ Analogue for task setting in the real world
 - squares represent tasks; chips represent resources; getting to task completion
 - vivid representation of resource
- ▶ Flexible form
 - manipulation of resources controlling
- ▶ Family of games that can differ in aspect

Perfect!!
Excellent!!



Social Preference Agent [Gal 06].

- ▶ Learns the extent to which people are affected by social preferences such as social welfare and competitiveness.
- ▶ Designed for one-shot leave-it scenarios
- ▶ Does not observe the outcome of its action

No previous data; too simple protocol

Y. Gal and A. Pfeffer. Predicting People's Bidding Behavior in Negotiation , AAMAS 2006.



Multi-Personality agent [TA05]

- ▶ Estimate the helpfulness and reliability of the opponents
- ▶ Adapt the personality of the agent accordingly
- ▶ Maintained Multiple Personality– one for each opponent
- ▶ Utility Function

S. Talman, Y. Gal, S. Kraus and M. Hadad. Adapting to Agents' Personalities in Negotiation, in AAMAS 2005.

CT Scenario

2

Agent &
human

- ▶ 4 ~~4~~ 5 players (all automated)
- ▶ Multiple rounds:
 - negotiation (flexible protocol),
 - chip exchange,
 - movements
- ▶ Incomplete information on others' chips
- ▶ Agreements are not enforceable
- ▶ Complex dependencies
- ▶ Game ends when one of the players:
 - reached goal
 - did not move for three movement phases.

Alternating
offers (2)

Complete
information

Summary of agents

- ▶ QOAgent
- ▶ KBAgent
- ▶ Gender-sensitive agent
- ▶ Social Preference Agent
- ▶ Multi-Personality agent

Personally, Utility, Rules Based agent (PURB)

Show PURB game

PURB: Cooperativeness

- ▶ **helpfulness trait:** willingness of negotiators to share resources
 - percentage of proposals in the game offering more chips to the other party than to the player
- ▶ **reliability trait:** degree to which negotiators kept their commitment
 - ratio between the number of chips

Build
cooperative
agent !!!



PURB: social utility function

- ▶ Weighted sum of PURB's and its partner's utility
- ▶ Person assumed to be using a truncated model (to avoid an infinite recursion):
 - The expected future score for PURB
 - *based on the likelihood that i can get to the goal*
 - The expected future score for nego partner
 - computed in the same way as for PURB
 - The cooperativeness measure of nego partner
 - *in terms of helpfulness and reliability,*
 - The cooperativeness measure of PURB by *nego partner*

PURB: Update of cooperativeness traits

- ▶ Each time an agreement was reached and transfers were made in the game, *PURB* updated both players' traits
 - values were aggregated over time using a discounting rate

Game 1

Board Display

		G		
me				

Player Chips Display

me	5	0	10	0	9
	1	0	13	10	0

Make Proposal- Pin ...

Propose to Responder:

Chips sent by ME

0	0	0	0	2
---	---	---	---	---

Chips sent by RESPONDER

0	0	0	2	0
---	---	---	---	---

Propose! Cancel

Make Proposal- Pin ...

Propose to Responder:

Chips sent by ME

0	0	0	0	1
---	---	---	---	---

Chips sent by RESPONDER

0	0	0	1	0
---	---	---	---	---


Propose! Cancel

Action History - Name 0

Initiator	Chips sent by ME	Chips sent by OTHER	Response
		1	transfer
me	0 0 0 0 1		transfer
	1	1	accept
		0 0 0 2 0	transfer
me	0 0 0 0 2		transfer
me	0 0 0 0 2	0 0 0 2 0	accept

Game 2

Board Display



Player Chips Display

me	5	0	10	0	9
○	1	0	13	10	0

Pin 300 Sent Propos...

Status: Pending Reply...

Proposer: me

Chips sent by Proposer: 0 0 0 1 0

Responder: ○

Chips sent by Responder: 2 0 0 0 0

Make Proposal- Pin ...

Propose to Responder: ○

Chips sent by ME: 0 0 0 0 1

Chips sent by RESPONDER: 0 0 0 1 0

Propose! Cancel

Action History - Name 0

Initiator	Chips sent by ME	Chips sent by OTHER	Response
me	0 0 0 0 1		transfer
me	0 0 0 0 1	0 0 2 0 0	accept
○		1	transfer
○	1	1	accept
○		0 0 0 2 0	transfer
me	0 0 0 0 2		transfer
me	0 0 0 0 2	0 0 0 2 0	accept

PURB's rules: utility function

- ▶ The weight of the negotiation partner's score in PURB's utility:
 - dependency relationships between participants: decreased when nego partner is independent
 - cooperativeness traits: increased with nego partner *cooperativeness measures*

PURB's rules principle

begins by acting reliably



Adapts over time to the individual measure of cooperativeness exhibited by its negotiation partner.

PURB's rules: *Accepting Proposals*

- ▶ Accepted an offer if its utility was higher than the utility from the offer it would make as a proposer in the same game state, or
- ▶ If accepting the offer was necessary to prevent the game from terminating


Game Board - Name 300

Game over!!! Your score is: 180

Phases Display

Phases	Time Left
Feedback Phase	0:01
> Communication Phase <	> 4:10 <
Exchange Phase	0:40
Movement Phase	0:20
Feedback Phase	0:01

Board Display



Player Chips Display

me	0	0	10	0	6
0	0	9	7	0	

PURB's rules: making proposals

- ▶ Generated a subset of possible offers
 - Cooperativeness traits of negotiation partner
 - dependency relationships
- ▶ Compute utility of the offers
- ▶ Non-deterministically chose any proposal out of the subset that provided a maximal benefit (within an epsilon interval).
- ▶ Examples:
 - if co-dependent and symmetric generate 1:1 offers
 - If PURB independent generate 1:2 offers

PURB's rules: Transferring Chips

- ▶ If the reliability of negotiation partner *was*
 - **Low:** do not send any of the promised chips.
 - **High:** send all of the promised chips.
 - **Medium:** the extent to which PURB was reliable depended on the dependency relationships in the game [randomization was used]
- ▶ Example: If *partner was task dependent, and the agreement makes it task independent, then PURB sent the largest set of chips such that partner remained task dependent.*

Experimental Design

- ▶ 2 conditions: Israeli (93) and U.S. (100)

- ▶ 3 types of instructions:
 - Movie of instruction;
 - Arabic instructions;
 - Instructions that are too complex; will not work well.

first offer

Co-dependent



Player Chips Display


me	1	0	13	10
○	5	0	19	0

9
0

Hypothesis

- ▶ People in the U.S. and Lebanon would differ significantly with respect to cooperativeness;
- ▶ An agent that modeled and adapted to the cooperativeness measures exhibited by people will play at least as well as people

Average Performance



Reliability Measures

	Co-dep	Task indep.	Task dep.	Average
People (Lebanon)	0.96	0.94	0.87	0.92

People (US)	0.64	0.78	0.51	0.65

Reliability Measures

	Co-dep	Task indep.	Task dep.	Average
PURB (Lebanon)	0.96	0.99	0.99	0.98

PURB (US)	0.59	0.59	0.72	0.62

Reliability Measures

	Co-dep	Task indep.	Task dep.	Average
PURB (Lebanon)	0.96	0.99	0.99	0.98
People (Lebanon)	0.96	0.94	0.87	0.92



PURB (US)	0.59	0.59	0.72	0.62
People (US)	0.64	0.78	0.51	0.65

Reliability Measures


	Co-dep	Task indep.	Task dep.	Average
PURB (Lebanon)	0.96	0.99	0.99	0.98
People (Lebanon)	0.96	0.94	0.87	0.92

PURB (US)	0.59	0.59	0.72	0.62
People (US)	0.64	0.78	0.51	0.65


Proposed offers vs accepted offers: average



Performance by Dependencies Lebanon




Performance by Dependencies U.S.




Co-dependent

No different in reaching the goal

Task Co-dependent (Lebanon)



Task Co-dependent (US)



Implications for agent design

- ▶ Adaptation to the behavioral traits exhibited by people lead proficient negotiation across cultures.
- ▶ In some cases, people may be able take advantage of adaptive agents by adopting ambiguous measures of behavior.

On going work Personality, Adaptive Learning (PAL) agent

- ▶ Data collected is used to build predictive models of human negotiation behavior:
 - Reliability
 - Acceptance of offers
 - Reaching the goal
- ▶ The utility function use the models
- ▶ Reduce the number of rules
- ▶ Limited search

G. Haim, Y. Gal and S. Kraus. Learning Human Negotiation Behavior Across Cultures, in HuCom2010.

Conclusions of PURB

- ▶ Presented a new agent-design that uses adaptation techniques to negotiate with people

Human-Computer
Negotiation: Learning
from Different
Cultures



proof of the benefit of the approach

Summary

**General opponent*
modeling improves agent
bargaining**


negotiations

Automated care-taker

I will be t
tired in
aftern


How can I
convince him?
What argument
should I give?

Let me try to
in



Specifications

How should I convince him to provide me with information?



Cycle: <10 hrs (patrol-S2-patrol); high Info retention: automated analysis

Which

Should I tell
him that we
are running
out of
antibiotics?

Build a game that
combines information
revelation and bargaining

Automated care-taker




I will be too tired in the afternoon!!!

I arrange for you to go to the physiotherapist in the afternoon

How can I convince him?
What argument should I give?




Security applications



Color Trails (CT)

- ▶ An infrastructure for agent design, implementation and evaluation for open environments
- ▶ Designed with Barbara Grosz (AAMAS 2004)
- ▶ Implemented by Harvard team and BIU team



Agents for Revelation Games

Peled Noam, Gal Kobi,
Kraus Sarit

Introduction - Revelation games

- Combine two types of interaction
- Signaling games (Spence 1974)
 - Players choose whether to convey private information to each other
- Bargaining games (Osborne and Rubinstein 1999)
 - Players engage in multiple negotiation rounds
- Example: Job interview

Colored Trails (CT)

Asymmetric

Board Display

G 50			○
? 50			
			? 50
me			? 50

Player Chips Display

me	0	10	0	14
○	0	0	10	14

Symmetric

Board Display

		G 50				me
? 50						
						? 50
○				? 50		

Player Chips Display

me	0	0	0	10	7	7
○	0	7	7	10	0	0

Why not equilibrium agents?

- ▶ Results from the social sciences suggest people do not follow equilibrium strategies:
 - Equilibrium based agents played against people failed.
- ▶ People rarely design agents to follow equilibrium strategies (Sarne et al AAMAS 2008).
- ▶ Equilibrium strategies are usually not cooperative – all lose.



Perfect Equilibrium (PE) Agent


- Solved using Backward induction.
- No signaling.
- Counter-proposal round (selfish):
 - Second proposer: Find the most beneficial proposal while the responder benefit remains positive.
 - Second responder: Accepts any proposal which gives it a positive benefit.

PE agent – Phase one


- First proposal round (generous):
 - First proposer: propose the opponent's counter-proposal.
 - First responder: Accepts any proposals which gives it the same or higher benefit from its counter-proposal.
- Revelation phase - revelation vs non revelation:
 - In both boards, the PE with goal revelation yields lower or equal expected utility than non-revelation

Benefits Diversity

Average proposed benefit to players from first and second rounds



Performance of PEQ agent



Revelation Effect


- Only 35% of the games played by humans included revelation
- Revelation had a significant effect on human performance but not on agent performance
- Revelation didn't help the agent
 - People were deterred by the strategic machine-generated proposals

	Revelation	No Revelation	Total
Human	114	89	98
Agent	62	54	58

SIGAL agent

Agent based on general opponent modeling:

Genetic
algorithm




Logistic
Regression


SIGAL Agent

- Learns from previous games.
- Predict the acceptance probability for each proposal using Logistic Regression.
- Models human as using a weighted utility function of:
 - Humans benefit
 - Benefits difference
 - Revelation decision
 - Benefits in previous round


Logistic Regression using a Genetic Algorithm




Comparing SIGAL learning vs Weka (Root mean squared error)




Comparing SIGAL learning vs Weka (correctly classified instances)



Expected benefit maximization



Maximization – round 2



Strategy Comparison

Strategies for the asymmetric board, non of the players has revealed, the human lacks 2 chips for reaching the goal, the agent lacks 1:

Agent	First round	Second round
PE	10 11 for 2	2 for 10 11
SIGAL	2 1 for 2	2 for 3 *


* In first round the agent was proposed a benefit of 90

Heuristics

- Tit for Tat
- Never give more than you asks in the counter-proposal
- Risk averseness
 - Isoelastic utility:

$$u(c) = \frac{c^{1-\rho}}{1-\rho}$$

$$\rho = 0.3$$



Learned Coefficients



- Responder benefit: $\uparrow$ (0.96)
- Benefits difference: $\downarrow$ (-0.79)
- Responder revelation: $\uparrow$ (0.26)
- Proposer revelation: $-$ (0.03)
- Responder benefit in first round: $\uparrow$ (0.45)
- Proposer benefit in first round: $\uparrow$ (0.33)

Methodology

- Cross validation.
 - 10-fold
- Over-fitting removal.
 - Stop learning in the minimum of the generalization error
- Error calculation on held out test set.
 - Using new human-human games
- Performance prediction criteria.

Performance

General opponent*
modeling improves
agent negotiations



General opponent* modeling in Maximization problems



AAT agent

Agent based on general* opponent modeling

Decision
Tree/
Naïve
Byes

AAT

Aspiration Adaptation Theory (AAT)


- ▶ Economic theory of people's behavior (Selten)
 - No utility function exists for decisions (!)
- ▶ Relative decisions used instead
- ▶ Retreat and urgency used for **goal variables**

Avi Rosenfeld and Sarit Kraus. Modeling Agents through Bounded Rationality Theories. Proc. of IJCAI 2009., JAAMAS, 2010.


Commodity search



Commodity search




Commodity search



If price < 800 buy; otherwise visit 5 stores and buy in the cheapest.

Results


Using AAT to Quickly Learn



Learning People's Negotiation Behavior

- ▶ KBAgent predicted the rank of the offered made by people
- ▶ Goal: predict the exact offers people made
- ▶ Problem: not enough data
- ▶ Approach: adding statistical knowledge of people's aspirations (AAT) to augment learning

Predicting People's Offers



» General
opponent*
modeling in
cooperative
environments

Coordination with limited communication

- ▶ Communication is not always possible:
 - High communication costs
 - Need to act undetected
 - Damaged communication devices
 - Language incompatibilities
 - Goal: Limited interruption of human activities

Zuckerman, S. Kraus and J. S. Rosenschein.
Using Focal Points Learning to Improve
Human-Machine Tactic Coordination, JAAMAS, 2010.


Focal Points (Examples)

- ▶ Divide £100 into two piles, if your piles are identical to your coordination partner, you get the £100. Otherwise, you get nothing.




101 equilibria

Focal points (Examples)




16 equilibria



9 equilibria

Focal Points

- ▶ Thomas Schelling (63)
- ▶ ***Focal Points*** = Prominent solutions to tactic coordination games



Prior work: Focal Points Based Coordination for closed environments

- ▶ Domain-independent rules that could be used by automated agents to identify focal points:
- ▶ Properties: Centrality, Firstness, Extremeness, Singularity.
 - Logic based model
 - Decision theory based model
- ▶ Algorithms for agents coordination

Kraus and Rosenchein MAAMA 1992


Fenster et al ICMAS 1995

Annals of Mathematics and Artificial Intelligence 2000

FPL agent

Agent based on general* opponent modeling

Decision
Tree/
neural
network




Focal Point


FPL agent

Agent based on general opponent modeling:

Decision
Tree/
neural
network




raw data
vector



FP vector

Focal Point Learning


- ▶ 3 experimental domains:



Results – cont'

General opponent*
modeling improves
agent coordination

- ▶ “very similar domain” (VSD) vs “similar domain” (SD) of the “pick the pile” game.



□ Random Guess □ Domain Data Agent □ FP Agent □ FPL Agent

**Experiments
with people is a
costly process**

Evaluation of agents (EDA)

- ▶ Peer Designed Agents (PDA): computer agents developed by humans
- ▶ Experiment: 300 human subjects, 50 PDAs, 3 EDA
- ▶ Results:
 - EDA outperformed PDAs in the same situations in which they outperformed people,
 - on average, EDA exhibited the same measure of generosity

R. Lin, S. Kraus, Y. Oshrat and Y. Gal. Facilitating the Evaluation of Automated Negotiators using Peer Designed Agents, in AAAI 2010.

Conclusions

- ▶ Negotiation and argumentation with people is required for many applications
- ▶ General* opponent modeling is beneficial
 - Machine learning
 - Behavioral model
 - Challenge: how to integrate machine learning and behavioral model

References

1. Y. Gal and A. Pfeffer Predicting People's Bidding Behavior in Negotiation , AAMAS 2006.
2. Y. Gal, S. Kraus, M. Gelfand, Y. Blumberg, and E. Salmon. An Adaptive Agent for Negotiating with People in Different Cultures, 2010.
3. C. M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *JAAMAS*, 15(2):221–252, 2007
4. S. Kraus and D. Lehmann. Designing and Building a Negotiating Automated Agent, *Computational Intelligence*, 11(1):132-171, 1995
5. S. Kraus, K. Sycara and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *AIJ*, 104(1-2):1-69, 1998.
6. R. Lin and Sarit Kraus. Can Automated Agents Proficiently Negotiate With Humans? *Communications of the ACM* Vol. 53 No. 1, Pages 78-88, January, 2010.
7. R. Lin, S. Kraus, Y. Oshrat and Y. Gal. Facilitating the Evaluation of Automated Negotiators using Peer Designed Agents, in *AAAI* 2010.
8. R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7):823–851, 2008
9. Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 2009
10. I. Zuckerman, S. Kraus and J. S. Rosenschein. Using Focal Points Learning to Improve Human-Machine Tactic Coordination, *JAAMAS*, 2010.

Tournament

- ▶ 2nd annual competition of state-of-the-art negotiating agents to be held in AAMAS'11

Do you want to participate?

At least \$2,000 for the winner!

Contact us!
sarit@cs.biu.ac.il