CWI
Centrum Wiskunde & Informatica

Artificial Intelligence
www.elsevier.com/locate/artint

NEC-AIST
AI Cooperative
Research Laboratory

NWO

Automated Negotiation League 2025

Tamara Florijn (CWI & Utrecht University)
Yasser Mohammad (NEC-AIST, Japan)
Tim Baarslag (CWI & Technical University Eindhoven)
ANAC board members

Motivating scenario

- Claire wants to plan two evenings to eat with a friend.
- Her two friends A and B live far apart while she lives in the middle, so she meets them separately.
- She first calls friend A to set a day, then calls friend B to set a day.

A

C

B

Preferences Claire

Monday	Tuesday	Wednesday	Utility
	V		0.7
V			0.5
		V	0.4
V	V		0.3
	V	V	0.3
V		V	0.9

What could she do?

Monday	Tuesday	Wednesday	Utility
	V		0.7
V			0.5
		V	0.4
V	V		0.3
	V	V	0.3
V		V	0.9

- She could call friend A, and propose Tuesday, since Tuesday is better then Monday or Wednesday!
- However, she cannot combine that deal with an appointment with friend B.

What could she do?

Monday	Tuesday	Wednesday	Utility
	V		0.7
V			0.5
		V	0.4
V	V		0.3
	V	V	0.3
V		V	0.9

- She could call friend A, and propose Tuesday, since Tuesday is better then Monday or Wednesday!
- However, she cannot combine that deal with an appointment with friend B.

Think ahead

Monday	Tuesday	Wednesday	Utility
	V		0.7
V			0.5
		V	0.4
V	V		0.3
	V	V	0.3
V		V	0.9

- The best option of all is one meeting on Monday, one on Wednesday.
- Monday has a higher utility, so she could propose Monday first to friend A.
- And after that, propose Wednesday to friend B.

This year's challenge

- Sequential negotiation
- Multi-deal negotiation
- In NegMAS

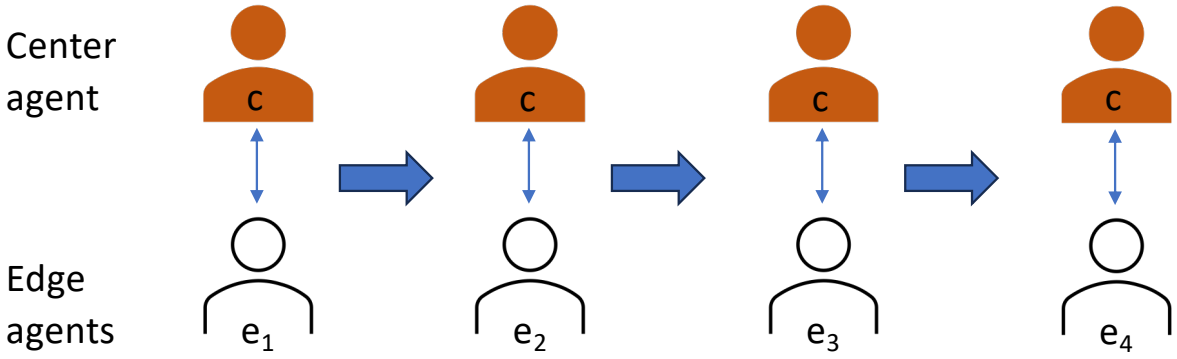

This year's challenge

- Unknown preferences
- Looking ahead over multiple negotiations
- Many options

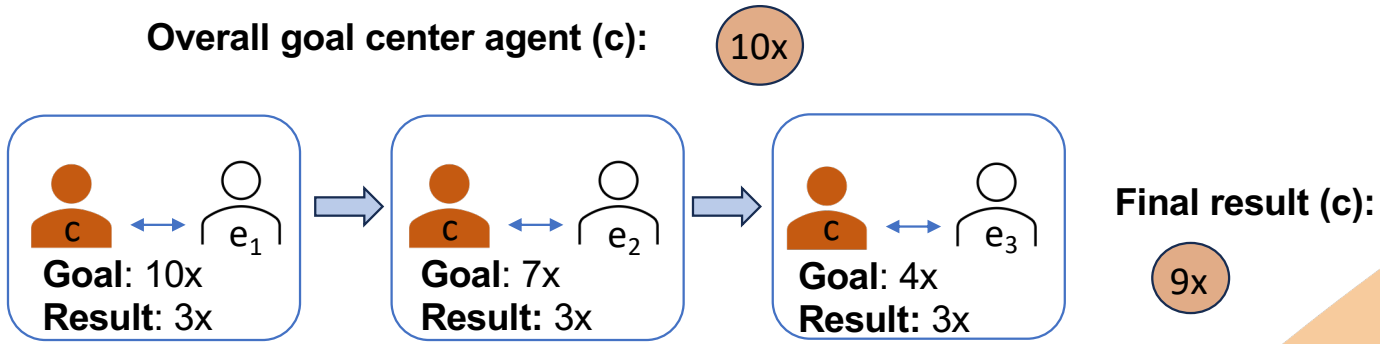

Evaluation

- Qualifications tournament
- Final tournament with top 12 of individual advantage.

Some numbers:

- Two independent finale runs (**1438** and **1453** repetitions).
- Every agent ran:
  - **21,570** and **21,795** times as center
  - **122,230** and **123,505** times as edge

Special thanks to Yasser Mohammad for his time and effort.

Evaluation

- Qualifications tournament
- Final tournament with top 12 of individual advantage.

- **Winners:**

  Top 3 individual advantage

$$A = \frac{\mu_c + \mu_e}{2}$$

Our participants

UfunATAgent

Smart Negotiator

OzUAgent

RUFL

- 17 groups of participants
- 7 countries

CARC2025

ProbaBot

RivAgent

SAC

TheMemorizer

StarGold

Analysis

Looking ahead

- Expected outcome (e.g RUFLagent, Probabot, RivAgent)
- Dynamic target (e.g. EOHAgent, CARCagent)
- Reinforcement learning techniques (e.g. SacAgent)

Analysis

Memory explosion

- The number of options grows quickly.
  - A small domain (100 outcomes) with 3 opponents: 1 million (10^6)
  - A medium domain (1000 outcomes) with 5 opponents: 1 quadrillion (10^{15})
- Sampling methods (e.g. The Memorizer, kAgent)
- Dynamic programming (e.g. Astrat3m)

Reveal of the winners

	Center	Edge	Score
1. ?			
1. ?			
3. ?			
4. A4e	0.6498	0.0874	0.3686
5. Wagent	0.6095	0.0897	0.3496
6. JeemNegotiator	0.5987	0.0845	0.3416
7. NayesianNiceAstrat3m	0.5941	0.0882	0.3412
8. ProbaBot	0.4958	0.0873	0.2916
9. SmartNegotiator	0.4898	0.0769	0.2833
10. KDY	0.4336	0.0910	0.2623
11. OzUAgent	0.4171	0.0881	0.2526
12. CARC2025	0.3628	0.0942	0.2285

Runner Up of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

?

Tamara Florijn Yasser Mohmmad

Reyhan Aydogan Katsuhide Fujita Tim Baarslag Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€100

Runner Up of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

This certificate is awarded to

Fukutoku Yuma

Tokyo University of Agriculture and Technology, Japan

to honor their excellent negotiation strategy

UfunATAgent

Tamara Florijn — Yasser Mohmmad

Reyhan Aydogan — Katsuhide Fujita — Tim Baarslag — Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€100

Winner of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

?

Tamara Florijn — Yasser Mohmmad

Reyhan Aydogan — Katsuhide Fujita — Tim Baarslag — Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€250

Winner of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

This certificate is awarded to

Hosssein Savari

University of Tehran, Iran

to honor their excellent negotiation strategy

SAC Agent

Tamara Florijn — Yasser Mohmmad

Reyhan Aydogan — Katsuhide Fujita — Tim Baarslag — Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€250

Winner of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

?

Tamara Florijn — Yasser Mohmmad

Reyhan Aydogan — Katsuhide Fujita — Tim Baarslag — Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€250

Winner of the Automated Negotiating Agents Competition 2025

~Automated Negotiation League~

(Individual Advantage)

This certificate is awarded to

Garrett Seo

Rutgers University, USA

to honor their excellent negotiation strategy

RUFL

Tamara Florijn — Yasser Mohmmad

Reyhan Aydogan — Katsuhide Fujita — Tim Baarslag — Catholijn Jonker

SPONSORED BY

NEC-AIST AI Cooperative Research Laboratory

Artificial Intelligence www.elsevier.com/locate/artint

CWI

NWO

€250

RUTGERS

Utility Fit Lookahead Agent – Sequential Multi-deal Negotiation

Work Done By: Garrett Seo, Tri-an Nguyen, Xintong Wang

Presented By: Garrett Seo

Sequential Multi-Deal Negotiation

- Center agent encounters multiple edge agents in sequence
- Subnegotiation (center <-> edge): Bilateral, Alternating Offers Protocol

- Rewarded for combination of all agreements
- Aware of own utility function
  - Opponent's utility function unknown

Baarslag, T. (2024). *Multi-deal Negotiation.* ANL 2025 Call for Participation.

A Look into Subnegotiations

- Let Ω denote the outcome space
  - Let Ω_i denote the i-th subnegotiation
  - $\Omega = \Omega_1 \times \cdots \times \Omega_n$
- Given utility function
  - $u : \Omega \to \mathbb{R}$
- What's the utility of realizing some suboutcome $\omega_i \in \Omega_i$?
  - At each subnegotation, what is best agreement?

- Naively,
  - We can underestimate the utility of ω_i
  - Let h be the history of previous suboutcomes
  - $u(\omega = (h, \omega_i, None, \ldots, None))$
- **Can we do better?**

Our Approach

Divide center strategy in two parts:

1. Lookahead planning
   - Utility estimation
   - Early termination
2. Conceding Strategy
   - Utility Fit
     - Opponent bids -> our estimated utility

Tree Representation To Solve Estimated Utilities

- Calculate estimated utilities of suboutcomes
- A node at depth i
  - Beginning of subnegotiation i
  - Contains suboutcomes from previous subnegotiations $0, \dots, i-1$
- Children represents all suboutcomes Ω_i
- Recursively,
  - Calculate expected utility of all children
  - Assign probabilities to children
    - Cooperative vs. adversarial
  - Propagate expected utility of parent upwards

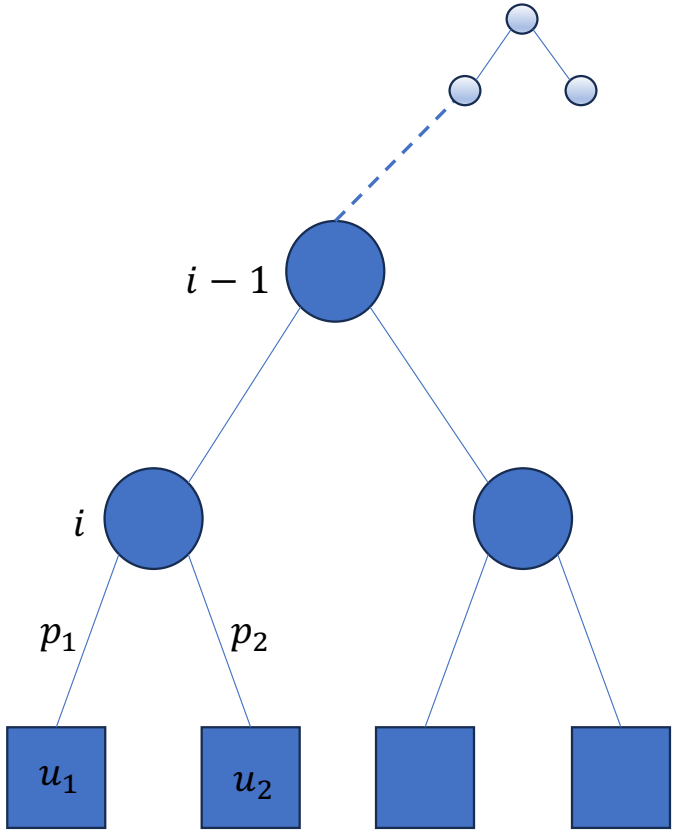

Early Termination

- What if the number of outcomes becomes intractable?
  - M suboutcomes, N subnegotiations → M^N outcomes
- Perform early termination
  - Do at each subnegotiation
- Stop at some depth k and propagate some terminal utility
  - Need heuristic
  - Use the naive underestimation
  - $u(\omega = (h, \omega_i, None, \dots, None))$

Concession Strategy – Time-Based

Assumption

- Opponents follow time-based strategy
- Faratin et al. (1998)

General Form

$$x(t) = u_{min} + (u_{max} - u_{min}) * (1 - t^{\mathrm{B}})$$

- $x(t)$: **opponent's** utility of bid offered at time t
- u_{min} : minimum utility
- u_{max} : maximum utility
- B : concession degree

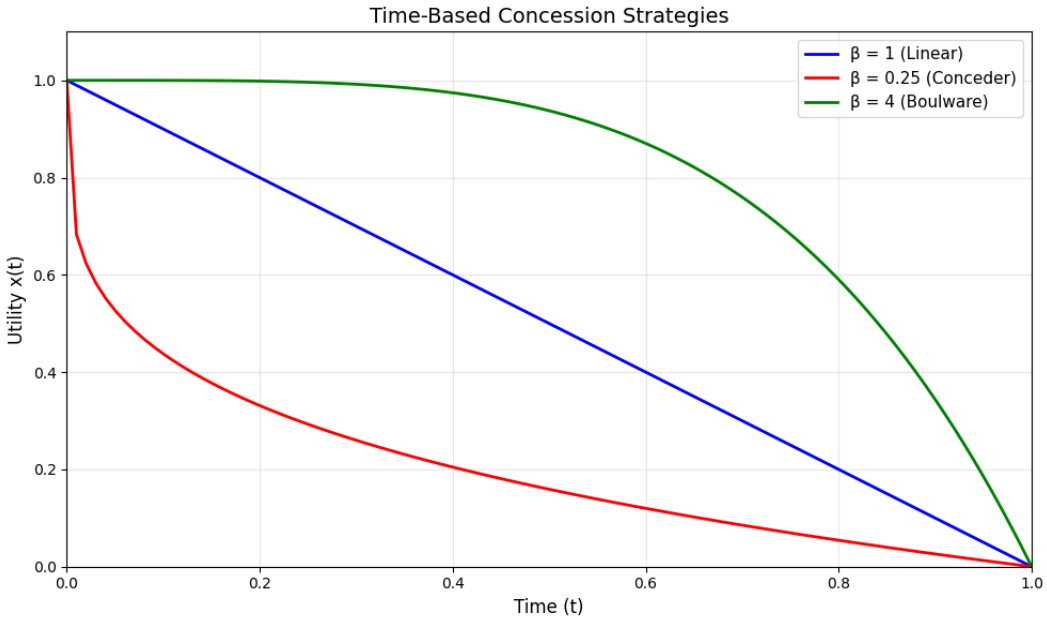

Our Concession Strategy

Idea:

- No discount
- Delay negotiation as long as possible
  - Reject all offers
  - Bid suboutcomes with large utility
- Make most informed bid at last timestep
  - From their offers

How To Use Information?

- As opponent concedes over time (adversarial)
  - Opponent utility decreases
  - Our utility increases
- Map opponent offers to our estimated utilities from lookahead
  - Estimated utilities increase w/ time
- Use opponent offers → our utility to fit utility curve

Utility Fit

$x'(t; \boldsymbol{u_{max}}, \mathbf{B})$

- x′ : **our estimated utility** from opponent's offers
- $\boldsymbol{u_{max}}$: our maximum utility, opponent willing to concede to
- $\mathbf{B}$: opponent's concessive degree

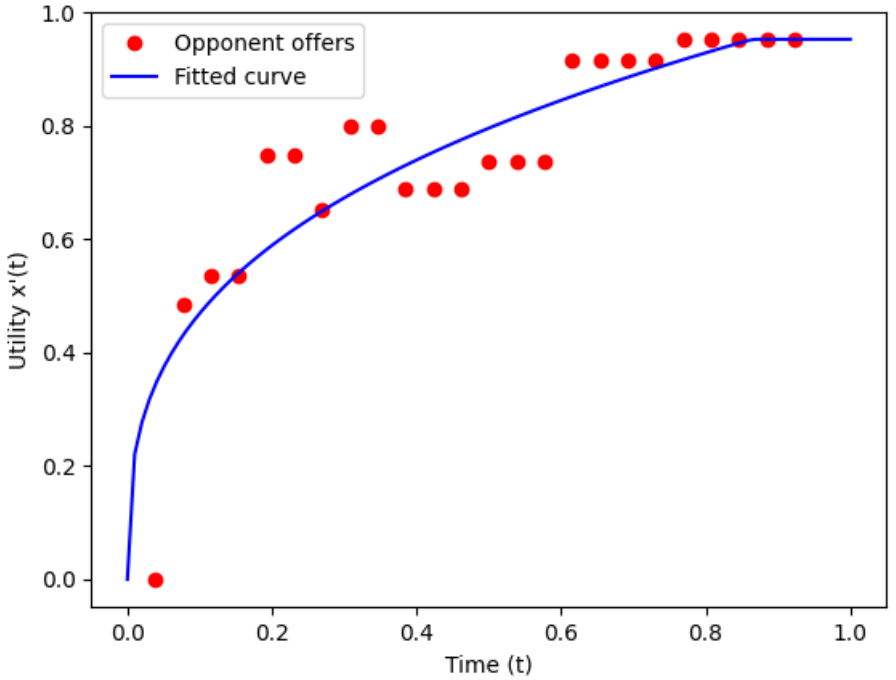

- Find parameters $\boldsymbol{u_{max}}$ and $\mathbf{B}$ that best fit according to their offers
- Final timestep:
  - Propose outcome w/ estimated utility: $x'(t = 1)$

CWI
Centrum Wiskunde & Informatica

Artificial Intelligence
www.elsevier.com/locate/artint

NEC-AIST
AI Cooperative
Research Laboratory

NWO

Automated Negotiation League 2025