

[COMP25] The Automated Negotiating Agents Competition (ANAC) 2025 Challenges and Results

Reyhan Aydoğan^{1,2}, Tim Baarslag^{3,4}, Tamara C.P. Florijn^{3,5}, Katsuhide Fujita^{6,7}, Catholijn M. Jonker^{2,8}, Yasser Mohammad⁷

¹Özyeğin University

²Delft University of Technology

³Centrum Wiskunde & Informatica (CWI)

⁴Eindhoven University of Technology

⁵Utrecht University

⁶Tokyo University of Agriculture and Technology

⁷NEC Corporation

⁸Leiden University

reyhan.aydogan@ozyegin.edu.tr, T.Baarslag@cw.nl, tamara.florijn@cw.nl, katfuji@cc.tuat.ac.jp,
C.M.Jonker@tudelft.nl, y.mohammad@nec.com

Abstract

This paper presents the primary research challenges and key findings from the 15th International Automated Negotiating Agents Competition (ANAC 2025), one of the official competitions of IJCAI 2025. We focus on two critical domains: multi-deal negotiations and the development of agents capable of concurrent negotiation within complex supply chain management environments. Furthermore, this work analyzes the results of the competition and outlines strategic directions for future iterations.

1 Introduction

Since 2010, the International Automated Negotiating Agents Competition (ANAC) has spearheaded innovation in autonomous agent research, introducing foundational challenges to the multi-agent systems community [Jonker *et al.*, 2017]. Since its inception with 7 participants [Baarslag *et al.*, 2012], the competition has expanded significantly, surpassing 850 total participants across 16 editions [Baarslag *et al.*, 2015]. The 2025 competition featured two distinct leagues. The Automated Negotiation League (ANL) required participants to develop bilateral agents that can engage in sequential multi-deal negotiations with multiple opponents and interdependent utility functions. In the Supply Chain Management League (SCML) [Mohammad *et al.*, 2019], the objective is to design agents that maximize profit within a competitive market. These agents must strategically navigate concurrent negotiations to secure raw materials and sell manufactured goods.

The 2025 edition attracted 142 international researchers forming 42 teams, competing for a 2,500 EUR prize pool. Both leagues utilized the NegMAS [Mohammad *et al.*, 2021] framework for agent development and tournament execution.

Agents competed in a round-robin format across randomized scenarios. ANL performance was evaluated via individual utility and Nash distance, whereas SCML rankings were determined by cumulative profit within the market simulation. The subsequent sections detail the specific league configurations and analyze the final results.

Moreover, ANAC 2025 introduced a pilot competition between IJCAI 2025 participants in which these participants negotiated against autonomous agents in a variety of scenarios for a 1,000 EUR prize pool. The scenarios included a trade negotiation inspired by the Supply Chain Management League, an Island survival tool distribution negotiation and a Grocery items distribution scenario. The participants conducted 1,456 negotiations and two winners were announced at the closing session of IJCAI 2025.

2 Main League Challenges & Results

In the main league of 2025, each participating agent encounters multiple opponents in sequence and is rewarded for the specific combination of the deals made in each negotiation. The challenge of this league is *sequential multi-deal* negotiation, where the participants need to coordinate an enormous number of options available and design an agent that outperforms other contestants by conceding efficiently and obtaining the best deals.

At the start of each negotiation, one agent will receive the role *center agent*, while the others receive the role *edge agent*. The center agent negotiates against multiple opponents (the edge agents) in a sequential manner, in what we call *subnegotiations*. Each subnegotiation will follow the Alternating Offers Protocol [Aydoğan *et al.*, 2014] where the starting agent makes an opening offer, which is followed by acceptance, a counteroffer, or a walk-away, repeated in a turn-taking fashion. If the negotiation ends, the center agent continues to the subsequent subnegotiation.

The agents negotiate over special types of domains that are

Rank	Agent	As Center	As Edge	Final Score
1	RUFL	0.714	0.084	0.399
1	SAC Agent	0.733	0.064	0.399
3	UFuncAtAgent	0.686	0.078	0.382

Table 1: ANL 2025 Results

characterized by the utility function of the center agent, two of which we discuss here. In the *Job Hunt scenario*, a job hunter (center agent) negotiates about two issues: the number of office days and the salary. In the end, the *center utility function* is defined in such a way that he receives the maximum of all the deals that he gathered. His preferences within a subnegotiation (called *side utility function*) are modeled using a linear additive function. Secondly, the *Target Quantity scenario* models a buyer (center agent) that negotiates with multiple sellers over the purchase quantity of a single product. The buyer aims for a specific quantity, for example to buy 10 products i.e. a peak preference at 10 products, while buying more or less gradually decreases the utility to 0.

In total, 17 teams submitted to ANL 2025, of which 12 were selected as finalists. The results of this year are shown in Table 1. As the negotiation protocol is sequential, the agents need to take into account what future actions will influence their current best decision. When the number of bids and thus the number of possible future actions grows, the number of combined outcomes to consider grows exponentially.

One of the winner agents, RUFL, tackled the computational challenge by constructing a tree to estimate the expected utility. The agent maintains a probability distribution over all the expected values of the child outcomes where their probabilities are generated via a softmax over their expected values with some temperature. By limiting the tree search at a certain level, they keep their computational time within limits. The second winner agent, SAC Agent, is built using the Soft Actor-Critic (SAC) reinforcement learning framework. The SAC bidding policy is a time-dependent concession strategy, where the concession rate is guided by the SAC model.

Other participants approach the uncertain future deals and their influence on the search space in distinct ways:

- **Pessimistic.** These agents pretend that no further deals will be made (e.g., UfuncAtAgent, SAC, default ANL 2025 strategies).
- **Contingent.** These agents use a probability distribution to assess the likelihood of further deals in the future (e.g., ProbaBot, RUFL, RivAgent, WAgent).
- **Optimistic.** These agents assume that the future will unfold exactly as planned (e.g., WAgent under conditionals).

Agents using the pessimistic methods circumvent the computational search challenge by ignoring future deals, disregarding any influence of future deals, limiting their performance. However, using contingent or optimistic approaches takes more computational time. All in all, the biggest challenge turned out to be the memory explosion caused by the big combined outcome spaces (e.g., CARC2025, OzUA-

gent, KDY, SmartNegotiator, ProbaBot), tackled with different techniques such as a dynamic target (e.g., EOHAgent, CARCAgent2025), sampling methods (e.g., the Memorizer, kAgent), dynamic programming (Astrat3m) or reinforcement learning techniques (e.g., SacAgent).

3 Supply-Chain Management League

The 2025 Supply Chain Management League (SCML) comprised two tracks: OneShot and Standard. A total of 20 qualified teams participated, with 11 competing in the OneShot track and 9 in the Standard track. Because the game mechanics remained identical to SCML 2024, participants were able to leverage the open-source code and technical reports of previous agents.

To further support the research community, the 2025 edition simplified reinforcement learning and MARL integration by providing a standard Gymnasium [Towers *et al.*, 2024] and a petting-zoo [Terry *et al.*, 2021] environments that encapsulate the SCML simulation [Mohammad *et al.*, 2025].

The primary objective of SCML is to align automated negotiation research with the practical challenges of industrial applications. The simulation serves as an abstraction of a fundamental procurement dilemma: *How can production needs be met while minimizing the combined costs of procurement and inventory?* In the SCML OneShot world, multiple autonomous agents manage factories within a supply chain, buying raw materials and selling final products. The market is driven by a set of exogenous un-negotiable contracts for the raw materials and final products. Every simulated days, agents must reach agreements with their suppliers/consumers for buying their production needs and selling their produced items. Contracts specify price and quantity (and delivery date in case of SCML-Standard); failure to meet agreed-upon obligations results in penalties. The ultimate goal for all agents is to maximize their accumulated profit.

SCML’s two tracks offer different levels of complexity regarding the negotiation search space and temporal dependencies:

- **OneShot Track:** Products are perishable, meaning profits or losses on a given day are independent of other days, except in the indirect effect of changing partners’ future negotiation behavior. Agents act as either buyers or sellers (never both), and the search space is limited by small price and quantity ranges. The research focus here is on *repeated concurrent negotiation* in many-to-many environments.
- **Standard Track:** This track introduces higher complexity by making products nonperishable (carrying storage costs) and allowing agents to negotiate delivery dates. The production graph is deeper, requiring agents to manage both buying and selling simultaneously. This shifts the challenge toward designing strategies for *dependent sequential sets of concurrent negotiations*.

The performance of the participants is summarized in Table 2. Notably, the winners of both tracks prioritized domain-specific heuristics over complex opponent modeling (for the second year in a row). Since 2020, all finalists of SCML

Track	Rank	Agent	Score
OnShot		CautiousOneShotAgent	1.0915
	1	CostAverseAgent	1.0896
	1	Rchan	1.0895
	1	AlmostEqualAgent	1.0892
Standard	1	AS0	1.009
		PenguinAgent	0.992
	2	XenoSotaAgent	0.960
	3	Ultra Super Miracle Final Agent Z	0.937

Table 2: SCML 2025 Results

could outperform the winner of the previous year. This pattern started to break in 2025 with the winner of SCML-OneShot 2024 (CautiousOneshotAgent) remaining the top performing agent in 2025 and the winner of SCML-Standard 2024 (PenguinAgent) only outperformed by the winner of SCML-Standard 2025 (AS0).

Three agents tied for the winner of SCML-Oneshot: 1) CostAverseAgent (developed by Yuzuru Kitamura) which focuses on minimizing the risk of financial loss by prioritizing cost-based decision-making over aggressive profit-seeking. 2) Rchan (developed by Shota Takayama) which is built around a sophisticated dynamic aspiration model and market estimation. 3) AlmostEqualAgent (developed by Kaito Miwa) which utilizes a strategy centered on proportional equity and collaborative surplus sharing to reach agreements quickly.

The winner of SCML-Standard (AS0 developed by Atsunaga et al.) focuses on a risk-balanced production-first strategy that adapts to the complexities of a full supply chain simulation. It employs a conservative negotiation strategy that sets price floors using expected manufacturing costs and storage fees, ensuring that every contract contributes to a positive net profit.

4 Conclusion

The 16th International Automated Negotiating Agents Competition (ANAC 2025) has once again demonstrated the evolving complexity of autonomous negotiation in multi-agent systems. By shifting the focus toward sequential multi-deal negotiations in ANL and deepened supply chain dependencies in the SCML, the competition continues to push the boundaries of how agents handle high-dimensional outcome spaces and temporal uncertainty.

Our analysis of the 2025 results yields several key insights:

- **Computational Scalability:** In ANL, the transition to multi-deal scenarios highlighted a critical trade-off between computational overhead and strategic foresight. Success was largely defined by an agent’s ability to manage “memory explosion” through tree-search pruning, sampling, or reinforcement learning (e.g., the SAC Agent).
- **Heuristics vs. Complexity:** In SCML, the continued success of domain-specific heuristics and risk-averse strategies (e.g., AS0 and CostAverseAgent) suggests

that in highly volatile, concurrent markets, robustness and cost-containment often outperform complex opponent modeling.

- **Human-Agent Interaction:** The 2025 pilot competition at IJCAI provided a rare benchmark for how autonomous agents perform against human negotiators in survival and trade scenarios, highlighting the need for more “explainable” and human-centric negotiation strategies.

Looking forward, the competition aims to further bridge the gap between theoretical frameworks and industrial application. Future iterations of ANAC will likely explore the integration of Large Language Models (LLMs) for more nuanced communication in the Human-Agent Negotiation League planned for 2026. ANL will explore the design of a negotiation agent for bilateral negotiation that tries to mislead its opponent. The agent is rewarded for the agreement made in the negotiation, as well as for how well it is able to deceive its opponent. This negotiation-in-the-wild challenge will explore the role of deception and its counter-measures in automated negotiation.

As usual with ANAC since 2010, all agent code is open-sourced and is accessible from the submission website at <https://anac.cs.brown.edu>.

As agent-based negotiation moves closer to real-world applications, the lessons learned from ANAC 2025 provide a vital roadmap for developing resilient, efficient, and collaborative autonomous systems.

Ethical Statement

The Human-Agent Negotiation pilot design received the approval of Ozyegin University, Turkey. Personal information of participants were not kept beyond the announcement of winners.

Acknowledgments

This publication is part of the Vidi project COMBINE (VI.Vidi.203.044) (partly) financed by the Dutch Research Council (NWO). ANAC 2025 was sponsored by AI Journal, NWO, and NEC-AIST AI Cooperative Research Laboratory.

References

- [Aydoğan et al., 2014] Reyhan Aydoğan, Koen V. Hindriks, and Catholijn M. Jonker. Multilateral mediated negotiation protocols with feedback. In *Novel Insights in Agent-based Complex Automated Negotiation*, pages 43–59. Springer Japan, Tokyo, 2014.
- [Baarslag et al., 2012] Tim Baarslag, Koen Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 113–135, Berlin, Heidelberg, 2012. Springer-Verlag.

- [Baarslag *et al.*, 2015] Tim Baarslag, Reyhan Aydoğan, Koen V. Hindriks, Katsuhide Fujita, Takayuki Ito, and Catholijn M. Jonker. The automated negotiating agents competition, 2010-2015. *AI Magazine*, 36(4):115–118, 12/2015 2015.
- [Jonker *et al.*, 2017] Catholijn M Jonker, Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen V Hindriks. Automated negotiating agents competition (anac). In *AAAI*, pages 5070–5072, 2017.
- [Mohammad *et al.*, 2019] Yasser Mohammad, Enrique Areyan Viqueira, Nahum Alvarez Ayerza, Amy Greenwald, Shinji Nakadai, and Satoshi Morinaga. Supply chain management world. In Matteo Baldoni, Mehdi Dastani, Beishui Liao, Yuko Sakurai, and Rym Zalila Wenkstern, editors, *PRIMA 2019: Principles and Practice of Multi-Agent Systems*, pages 153–169, Cham, 2019. Springer International Publishing.
- [Mohammad *et al.*, 2021] Yasser Mohammad, Shinji Nakadai, and Amy Greenwald. Negmas: a platform for situated negotiations. In *Recent Advances in Agent-based Negotiation: Formal Models and Human Aspects*, pages 57–75. Springer, 2021.
- [Mohammad *et al.*, 2025] Yasser Mohammad, Shinji Nakadai, and Amy Greenwald. Automated negotiation in supply chains a generalist environment for rl/marl research. In Ryuta Arisaka, Victor Sanchez-Anguix, Sebastian Stein, Reyhan Aydoğan, Leon van der Torre, and Takayuki Ito, editors, *PRIMA 2024: Principles and Practice of Multi-Agent Systems*, pages 19–24, Cham, 2025. Springer Nature Switzerland.
- [Terry *et al.*, 2021] J Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. Pettingzoo: Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34:15032–15043, 2021.
- [Towers *et al.*, 2024] Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, et al. Gymnasium: A standard interface for reinforcement learning environments. *arXiv preprint arXiv:2407.17032*, 2024.