Transfer Learning based Agent for Automated Negotiation

Extended Abstract

Siqi Chen*
College of Intelligence and Computing, Tianjin University
Tianjin, China
siqichen@tju.edu.cn

Qisong Sun
College of Intelligence and Computing, Tianjin University
Tianjin, China
qisong_sun@tju.edu.cn

Heng You
College of Intelligence and Computing, Tianjin University
Tianjin, China
hengyou@tju.edu.cn

Tianpei Yang
University of Alberta
Edmonton, Canada
tpyang@tju.edu.cn

Jianye Hao
College of Intelligence and Computing, Tianjin University
Tianjin, China
jianye.hao@tju.edu.cn

ABSTRACT

Although great success has been made in automated negotiation, a major issue still stands out: it is inefficient that learning a policy from scratch when an agent encounters an unknown opponent. Transfer learning (TL) can alleviate this problem by utilizing the knowledge of previously learned policies to accelerate the current task learning. This work presents a novel Transfer Learning-based Negotiating Agent (TLNAgent) framework that allows an autonomous agent to transfer previous knowledge from source policies to help with new tasks, while boosting its performance. TLNAgent comprises three key components: the negotiation module, the adaptation module and the transfer module. Specifically, the negotiation module is responsible for interacting with the other agent during negotiation. The adaptation module measures the helpfulness of each source policy based on a fusion of two selection mechanisms. The transfer module is based on lateral connections between source and target networks and accelerates the agent's training by transferring knowledge from the selected source policy. Our comprehensive experiments clearly demonstrate that TL is effective in the context of automated negotiation, and TLNAgent outperforms state-of-the-art negotiating agents in various domains.

KEYWORDS

Automated negotiation; Agreement Technologies; Transfer learning; Reinforcement learning; Deep learning

ACM Reference Format:

Siqi Chen, Qisong Sun, Heng You, Tianpei Yang, and Jianye Hao. 2023. Transfer Learning based Agent for Automated Negotiation: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 4 pages.

1 INTRODUCTION

In automated negotiation, autonomous agents attempt to reach a joint agreement on behalf of human negotiators in a buyer-seller or consumer-provider setup [13, 27, 42]. The biggest driving force behind research into automated negotiation is arguably the augmentation of the abilities of human negotiators as well as the broad spectrum of potential applications in industrial and commercial domains [e.g., 12, 18, 33, 44]. The interaction framework enforced in automated negotiation lends itself to the use of machine learning techniques for exploring effective strategies. Inspired by advances in deep learning [8, 17, 22, 36] and reinforcement learning (RL) [16, 20, 39, 46], the application of deep RL on negotiation has made significant success [2, 4, 7, 8, 11, 25, 43]. However, all these methods need to learn from scratch when faced with new opponents, which is inefficient and impractical.

The existing works mainly focus on how to use the gained experience to train an agent to deal with the encountered opponents [1, 2, 5, 6, 9, 25, 35, 43]. In practice, the agent however may be faced with unfamiliar or unknown opponent strategies, in which its policy may be ineffective, and the agent thus needs to learn a new policy from scratch [21, 23, 30]. Besides, in most negotiation settings, agents are required to negotiate with multiple types of opponents in turn which may be unknown [3, 19, 31, 41]. The problem behind it is that learning in such a manner is time-costly and may also restrict its potential performance (e.g., ignoring all previous experience and learned policies that are relevant to the current task). So, a core question arises: how to accelerate the learning process of new opponent strategy, while improving the performance of the learned policy.

This paper describes an attempt to answer the question with transfer learning (TL), which has emerged as a promising technique to accelerate the learning process of the target task by leveraging prior knowledge [10, 28, 32, 38, 48]. We propose a novel TL-based negotiating agent called TLNAgent, which is the first framework to apply TL in automated negotiation. It comprises three key components: the negotiation module, the adaptation module, and the transfer module. The negotiation module is responsible for interacting with other agents in the negotiation and providing information for other modules. The adaptation module measures the helpfulness of the source task concurrently based on two metrics: similarity between the source opponents and the current opponent, as well as the specific performance of the source policies on the target task [14, 24, 34, 47]. The transfer module is the core of our agent

*Corresponding author: Siqi Chen

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

framework, which accelerates the agent's training utilizing the source policies that the adaptation module selects. The comprehensive experiments conducted in the work clearly demonstrate the effectiveness of TLNAgent.

2 TRANSFER LEARNING BASED AGENT

To enable the agent to reuse the learned knowledge and learn how to deal with new opponents, we firstly propose the **Transfer Learning Based Agent For Automated Negotiation** framework (See Figure 1). The framework is composed of three modules: negotiation module, adaptation module, and transfer module. Through the cooperation of three modules, the framework can accelerate the learning process when encountering a new opponent and improve the learned policy performance [15, 26, 37, 45]. Our framework performs much better than traditional methods based on RL, which will be validated in our experiments.

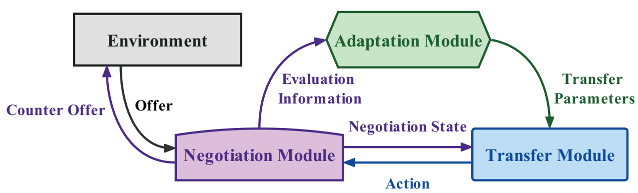

Figure 1: An overview of our framework

Negotiation Module is used to interact with other negotiating agents (e.g., receiving offers from opponents, generating counter-offers and making acceptance/rejection decisions). It also provides the necessary information for the adaptation and transfer module.

Adaptation Module decides when and which source policies are more appropriate to be transferred in the current task. To measure the transferability of each source policy, we propose two evaluation metrics: (a) **performance metric**, which represents the specific performance of the source policy on the target task, (b) **similarity metric**, which measures the similarity between the source opponents and the current opponent. Both evaluation metrics need the negotiation module to provide the necessary information. Subsequently, weighting factors resulting from the evaluation are passed to the transfer module.

Transfer Module is used to accelerate the learning process and boost performance encountering new opponents. After the adaptation module generates the weight factors, the transfer module extracts useful knowledge from source policies based on lateral connections [29, 40, 47], and then makes decisions for the negotiation module to obtain feedback. In this way, the transfer module allows our agent to leverage useful knowledge to learn a high-performing policy in the current environment.

3 EXPERIMENTS

Environments: To verify the efficient learning ability of TLNAgent for previously unknown opponents, we evaluate the agent with multiple tasks consisting of different opponents and domains. The following two transfer metrics are used in experiments:

(1) Time to threshold benchmark: the learning time TLNAgent and baselines required to achieve the convergence performance in a negotiation, which is denoted by the ratio of the convergence episode number and the total episode number;
(2) Transfer ratio: the ratio of mean utility obtained by the agent negotiating with a certain opponent over all 18 domains between TLNAgent and the learn from scratch baseline.

Baselines: To demonstrate the advantages of using previous knowledge and the superiority of the transfer method when faced with new opponents, we consider the following two baselines:

- Learn from scratch, which uses the standard DRL algorithm SAC and learns without prior knowledge in the new negotiation environment;
- Learn from teachers, which is directly trained by the opponents that are used to train the source policies.

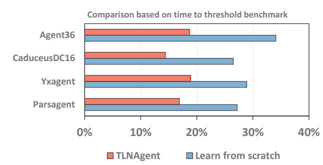

(a) Time to threshold benchmark

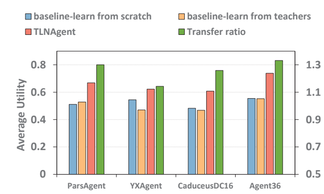

(b) Transfer ratio benchmark

Figure 2: Performance of TLNAgent and Learn from scratch on time to threshold benchmark and transfer ratio benchmark.

Figure 2a compares results of TLNAgent and other baselines on the time to threshold benchmark. As we can see, the TLNAgent converges faster than Learn from scratch in the face of different opponents. This means that the transfer module accelerates the agent's training utilizing the source policy that the adaptation module selects. As shown in Figure 2b, TLNAgent performs better for all opponents, achieving a 26% improvement in average utility compared to the two baselines. This is because TLNAgent transfers helpful knowledge from multiple source policies to the target task learning process through the transfer module.

4 CONCLUSION AND FUTURE WORK

In this paper, we introduced a novel transfer learning based negotiating agent framework called TLNAgent for effective and efficient automated negotiation. The framework contains three components: the negotiation module, the adaptation module and the transfer module. Furthermore, the framework adopts the performance metric and the similarity metric to measure the transferability of the source policies. The experimental results show a clear performance advantage of TLNAgent over available state-of-the-art agents (chosen from previous editions of ANAC competitions) in various aspects. TLNAgent opens several new research avenues, among which we consider the following as the most promising. First, as opponent modeling is another helpful way to improve the efficiency of a negotiation, it's worthwhile investigating how to combine opponent modeling techniques with our framework. Also, it is very interesting to see how well TLNAgent performs against human negotiators. The third important avenue we see is to enlarge the scope of the proposed framework to concurrent negotiations.

ACKNOWLEDGMENTS

This study is supported by the National Natural Science Foundation of China (Grant No. 61602391). Moreover, we greatly appreciate the valuable comments from Gerhard Weiss at Maastricht University to improve the quality of this paper.

REFERENCES

[1] Tim Baarslag, Koen V. Hindriks, Mark Hendrikx, Alexander Dirkzwager, and Catholijn M. Jonker. 2014. Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies. In *Novel Insights in Agent-based Complex Automated Negotiation*. 61–83.

[2] Pallavi Bagga, Nicola Paoletti, Bedour Alrayes, and Kostas Stathis. 2020. A Deep Reinforcement Learning Approach to Concurrent Bilateral Negotiation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*. 297–303.

[3] David Carmel and Shaul Markovitch. 1995. Opponent Modeling in Multi-Agent Systems. In *Adaption and Learning in Multi-Agent Systems, IJCAI'95 Workshop, Montréal, Canada, August 21, 1995, Proceedings (Lecture Notes in Computer Science, Vol. 1042)*, Gerhard Weiß and Sandip Sen (Eds.). Springer, 40–52.

[4] Ho-Chun Herbert Chang. 2021. Multi-issue negotiation with deep reinforcement learning. (2021).

[5] Siqi Chen, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss. 2013. Optimizing Complex Automated Negotiation using Sparse Pseudo-input Gaussian processes. In *Proceedings of the 12th Int. Joint Conf. on Automomous Agents and Multi-Agent Systems*. ACM, Saint Paul, Minnesota, USA, 707–714.

[6] Siqi Chen, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss. 2013. Using conditional restricted Boltzmann machine for highly competitive negotiation tasks. In *Proceedings of the 23th Int. Joint Conf. on Artificial Intelligence*. AAAI Press, 69–75.

[7] Siqi Chen and Ran Su. 2022. An autonomous agent for negotiation with multiple communication channels using parametrized deep Q-network. *Mathematical Biosciences and Engineering* 19, 8 (2022), 7933–7951.

[8] Siqi Chen, Qisong Sun, and Ran Su. 2022. An Intelligent Chatbot for Negotiation Dialogues. In *Proceedings of IEEE 20th International Conference on Ubiquitous Intelligence and Computing (UIC)*. IEEE, 68–73.

[9] Siqi Chen and Gerhard Weiss. 2014. An Intelligent Agent for Bilateral Negotiation with Unknown Opponents in Continuous-Time Domains. *ACM Trans. Auton. Adapt. Syst.* 9, 3 (2014), 16:1–16:24.

[10] Siqi Chen and Gerhard Weiss. 2015. An approach to complex agent-based negotiations via effectively modeling unknown opponents. *Expert Syst. Appl.* 42, 5 (2015), 2287–2304.

[11] Siqi Chen, Yang Yang, and Ran Su. 2022. Deep reinforcement learning with emergent communication for coalitional negotiation games. *Mathematical Biosciences and Engineering* 19, 5 (2022), 4592–4609.

[12] Lei Duan, Mustafa K. Dogru, Ulas Ozen, and J.Christopher Beck. 2012. A negotiation framework for linked combinatorial optimization problems. (2012), 158–182.

[13] Shaheen Fatima, Sarit Kraus, and Michael J. Wooldridge. 2014. *Principles of Automated Negotiation*. Cambridge University Press.

[14] Fernando Fernández and Manuela M. Veloso. 2006. Probabilistic policy reuse in a reinforcement learning agent. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*. 720–727.

[15] Katsuhide Fujita. 2014. Automated Negotiating Agent with Strategy Adaptation for Multi-times Negotiations. In *Recent Advances in Agent-based Complex Automated Negotiation [revised and extended papers from the 7th International Workshop on Agent-based Complex Automated Negotiation, ACAN 2014, Paris, France, May 2014] (Studies in Computational Intelligence, Vol. 638)*, Naoki Fukuta, Takayuki Ito, Minjie Zhang, Katsuhide Fujita, and Valentin Robu (Eds.). Springer, 21–37.

[16] Robert C. Grande, Thomas J. Walsh, and Jonathan P. How. 2014. Sample Efficient Reinforcement Learning with Gaussian Processes. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014 (JMLR Workshop and Conference Proceedings, Vol. 32)*. JMLR.org, 1332–1340.

[17] Shixiang Gu, Timothy P. Lillicrap, Ilya Sutskever, and Sergey Levine. 2016. Continuous Deep Q-Learning with Model-based Acceleration. In *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016 (JMLR Workshop and Conference Proceedings, Vol. 48)*, Maria-Florina Balcan and Kilian Q. Weinberger (Eds.). JMLR.org, 2829–2838.

[18] He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling Strategy and Generation in Negotiation Dialogues. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*. 2333–2343.

[19] Pablo Hernandez-Leal, Enrique Munoz de Cote, and Luis Enrique Sucar. 2015. Opponent Modeling against Non-stationary Strategies: (Doctoral Consortium). In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind (Eds.). ACM, 1989–1990.

[20] Shahin Jabbari, Matthew Joseph, Michael J. Kearns, Jamie Morgenstern, and Aaron Roth. 2017. Fairness in Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017 (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 1617–1626.

[21] Hamid Jazayeriy, Masrah Azrifah Azmi Murad, Md Nasir Sulaiman, and Nur Izura Udzir. 2011. The Learning of an Opponent's Approximate Preferences in Bilateral Automated Negotiation. *J. Theor. Appl. Electron. Commer. Res.* 6, 3 (2011), 65–84.

[22] Peter H. Jin, Kurt Keutzer, and Sergey Levine. 2018. Regret Minimization for Partially Observable Deep Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018 (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer G. Dy and Andreas Krause (Eds.). PMLR, 2347–2356.

[23] Haralambie Leahu, Michael Kaisers, and Tim Baarslag. 2019. Preference Learning in Automated Negotiation Using Gaussian Uncertainty Models. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 2087–2089.

[24] Siyuan Li and Chongjie Zhang. 2018. An Optimal Online Method of Selecting Source Policies for Reinforcement Learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*. 3562–3570.

[25] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2016. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

[26] Iou-Jen Liu, Jian Peng, and Alexander G. Schwing. 2019. Knowledge Flow: Improve Upon Your Teachers. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.

[27] Fernando Lopes, Michael Wooldridge, and A. Novais. 2008. Negotiation among autonomous computational agents: principles, analysis and challenges. (2008), 1–44.

[28] Emilio Parisotto, Lei Jimmy Ba, and Ruslan Salakhutdinov. 2016. Actor-Mimic: Deep Multitask and Transfer Reinforcement Learning. In *Proceedings of the 4th International Conference on Learning Representations*.

[29] Aaditya Ramdas, Nicolás García Trillos, and Marco Cuturi. 2017. On Wasserstein Two-Sample Testing and Related Families of Nonparametric Tests. (2017), 47.

[30] Hsin Rau, Mou-Hsing Tsai, Chao-Wen Chen, and Wei-Jung Shiang. 2006. Learning-based automated negotiation between shipper and forwarder. *Comput. Ind. Eng.* 51, 3 (2006), 464–481.

[31] Tjitze Rienstra, Matthias Thimm, and Nir Oren. 2013. Opponent Models with Uncertainty for Strategic Argumentation. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, Francesca Rossi (Ed.). IJCAI/AAAI, 332–338.

[32] Andrei A. Rusu, Sergio Gomez Colmenarejo, Çaglar Gülçehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. 2016. Policy Distillation. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.

[33] Victor Sanchez-Anguix, Vicente Julian, Vicente Botti, and Ana García-Fornes. 2013. Tasks for agent-based negotiation teams: Analysis, review, and challenges. (2013), 2480–2494.

[34] Simon Schmitt, Jonathan J. Hudson, Augustin Zídek, Simon Osindero, Carl Doersch, Wojciech M. Czarnecki, Joel Z. Leibo, Heinrich Küttler, Andrew Zisserman, Karen Simonyan, and S. M. Ali Eslami. 2018. Kickstarting Deep Reinforcement Learning. (2018).

[35] Ayan Sengupta, Shinji Nakadai, and Yasser Mohammad. 2022. Transfer Learning Based Adaptive Automated Negotiating Agent Framework. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*. 468–474.

[36] Ran Su, Haitang Yang, Leyi Wei, Siqi Chen, and Quan Zou. 2022. A multi-label learning model for predicting drug-induced pathology in multi-organ based on toxicogenomics data. *PLoS Computational Biology* 18, 9 (2022), e1010402.

[37] Yunzhe Tao, Sahika Genc, Jonathan Chung, Tao Sun, and Sunil Mallya. 2021. REPAINT: Knowledge Transfer in Deep Reinforcement Learning. In *Proceedings of the 38th International Conference on Machine Learning*. 10141–10152.

[38] Matthew E. Taylor and Peter Stone. 2009. Transfer Learning for Reinforcement Learning Domains: A Survey. (2009).

[39] Oriol Vinyals, Igor Babuschkin, and Wojciech M. Czarnecki et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. (2019), 350–354.

[40] Michael Wan, Tanmay Gangwani, and Jian Peng. 2020. Mutual Information Based Knowledge Transfer Under State-Action Dimension Mismatch. In *Proceedings of the Thirty-Sixth Conference on Uncertainty in Artificial Intelligence*. 1218–1227.

[41] Yuchen Wang, Fenghui Ren, and Minjie Zhang. 2017. Opponent Modeling with Information Adaptation (OMIA) in Automated Negotiations. In *Autonomous Agents and Multiagent Systems - AAMAS 2017 Workshops, Best Papers, São Paulo, Brazil, May 8-12, 2017, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 10642)*, Gita Sukthankar and Juan A. Rodríguez-Aguilar (Eds.). Springer, 21–35.

[42] G. Weiss. 2013. *Multiagent Systems, 2nd edition.*

[43] Leling Wu, Siqi Chen, Xiaoyang Gao, Yan Zheng, and Jianye Hao. 2021. Detecting and Learning Against Unknown Opponents for Automated Negotiations. In *PRICAI 2021: Trends in Artificial Intelligence*, Duc Nghia Pham, Thanaruk Theeramunkong, Guido Governatori, and Fenrong Liu (Eds.). 17–31.

[44] Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. 2021. Improving Dialog Systems for Negotiation with Personality Modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*. 681–693.

[45] Tianpei Yang, Jianye Hao, Zhaopeng Meng, Zongzhang Zhang, Yujing Hu, Yingfeng Chen, Changjie Fan, Weixun Wang, Wulong Liu, Zhaodong Wang, and Jiajie Peng. 2020. Efficient Deep Reinforcement Learning via Adaptive Policy Transfer. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*. 3094–3100.

[46] Deheng Ye and Guibin Chen et al. 2020. Towards Playing Full MOBA Games with Deep Reinforcement Learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems.*

[47] Heng You, Tianpei Yang, Yan Zheng, Jianye Hao, and Matthew E. Taylor. 2022. Cross-domain adaptive transfer reinforcement learning based on state-action correspondence. In *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*. 2299–2309.

[48] Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. 2020. Transfer Learning in Deep Reinforcement Learning: A Survey. (2020).