

AI for Mechanism Design and Strategic Decision Making (AIMS)

ICLR 2026 Workshop Proposal

1. Workshop Summary

Mechanism Design (MD) [1, 2, 3, 4] and Strategic Decision Making (SDM) [5, 6, 7] are foundational pillars of economic theory with growing relevance to computer science and multi-agent systems. MD focuses on designing rules and incentives to align the behaviors of self-interested agents with socially optimal outcomes. SDM, conversely, concerns the intelligent, adaptive processes agents use to make decisions under constraints and uncertainty. **MD and SDM are profoundly interconnected**: mechanism designers must anticipate strategic behaviors, while agents continuously adapt their strategies in response to mechanisms and environmental dynamics. Theoretically, they are integral to microeconomics, game theory, and computer science, with significant applications in AI. Industrially, MD and SDM tools effectively model large-scale problems in online advertising [8, 9], financial markets [10, 11, 12], e-commerce [13, 14], and energy trading [15, 16, 17].

Recent advances in AI [18, 19, 20, 21] are introducing novel methodologies that enhance and, in some cases, redefine MD and SDM. While conventional approaches often prioritize theoretical tractability through simplifying assumptions [22, 3, 23], modern data-driven techniques offer greater scalability in complex, high-dimensional environments. For instance, multi-agent reinforcement learning (MARL) provides robust frameworks for modeling complex SDM, simulating how agents learn and adapt to reveal emergent collective phenomena and equilibrium dynamics [24, 25, 26]. The emerging field of “differentiable economics” [27, 28, 29, 30, 31, 32, 33, 34] uses gradient-based optimization for end-to-end mechanism design, leveraging data to discover efficient solutions.

More recently, **foundation models and generative AI**, including Large Language Models (LLMs) [20, 35], have spurred a new wave of interdisciplinary research. These models’ superior language capabilities are broadening MD and SDM to more practical scenarios [36, 37, 38, 39]. Simultaneously, they are spawning novel research questions, such as analyzing LLMs’ behavior as quasi-rational agents [40, 41, 42, 43] and investigating output manipulation through auctions [44, 45, 46, 47, 48].

To capitalize on this emerging synergy, we propose a full-day workshop at ICLR 2026. This workshop will serve as a premier venue to **catalyze** interdisciplinary research at the intersection of modern AI, mechanism design, and strategic decision making. We aim to **bridge** distinct research communities by bringing together experts from machine learning, economics, theoretical computer science, and operations research. The program will **explore** how advanced AI can redefine core problems in MD and SDM and **showcase** impactful real-world applications and industrial case studies.

Call for Papers

We invite submissions on how modern AI can **redefine, extend, or automate** core problems in Mechanism Design (MD) and Strategic Decision Making (SDM). We strongly encourage submissions from diverse fields, including machine learning, reinforcement learning, multi-agent systems, optimization, NLP, human-AI interaction, economics, operations research, and theoretical computer science.

Topics of interest include, but are not limited to:

- **AI for Mechanism Design:**
 - Discovering optimal, robust, and adaptive mechanisms for auctions, matching, and voting.
 - Automating the validation of theoretical properties in novel mechanisms.
 - Extending classical mechanisms to high-dimensional or natural language settings.
 - Enhancing the interpretability and transparency of complex mechanisms.
- **AI for Strategic Decision Making:**
 - Assisting strategic decisions in complex, dynamic environments.
 - Characterizing and influencing dynamics and equilibria in multi-agent systems.
 - Improving human-machine collaboration in strategic settings.
- **Theory, Ethics, and Societal Impact:**
 - Formal models of AI agent interaction and evolution.
 - Theoretical analyses of benefits and pitfalls (e.g., efficiency, fairness, robustness).
 - Policy and governance frameworks for AI in MD/SDM.
 - Forecasting the societal impact of AI-assisted mechanisms and strategic interactions.
- **Applications and Case Studies:**
 - Real-world deployments and empirical insights (e.g., advertising, cloud markets).
 - Practical challenges in scaling AI for MD/SDM.
 - Novel benchmarks, datasets, and simulation platforms.
 - Blueprints for integrating AI into existing MD/SDM systems.

Submission Tracks: Submissions will be managed via OpenReview. Accepted papers will have their track indicated on the workshop website.

- **Long Paper Track:** Up to 9 pages + references for novel theoretical results, algorithms, empirical studies, comprehensive surveys, or real-world applications.
- **Short Paper Track:** Up to 4 pages + references for unpublished ideas, modest theoretical results, follow-up experiments, or fresh perspectives on existing work.

Review Process: Each submission will receive at least three reviews from our diverse Program Committee. Appendices are permitted but not guaranteed to be reviewed.

Publication and Presentation: All accepted papers are **non-archival** and will be featured on our website. A select number of top-rated **Long Papers** will be chosen for **oral presentations**. All other papers will be presented in a **poster session**. We will present **best paper awards** and offer **best poster awards** based on attendee voting.

We welcome cutting-edge research from academia and practical contributions from industry, including applied solutions, case studies, and real-world experiences (both successful and unsuccessful).

Comparison to Related Workshops

Several existing workshops explore the application of Machine Learning techniques within specific domains. Examples include those focused on online advertising (e.g., [AdKDD](#), [AI Driven Online Advertising](#), [Frontiers of Online Advertising](#)) and e-commerce (e.g., [LLMs for E-Commerce](#), [Decision Intelligence and Analytics for Online Marketplaces](#)). While the topics in these workshops may intersect with MD and SDM, their primary emphasis is typically on applied models tailored to those specific contexts, such as click-through rate prediction in online advertising or product valuation in e-commerce. Our workshop, in contrast, will maintain a more concentrated focus on the foundational aspects of MD and SDM with emerging AI techniques.

Another stream of workshops investigates the role of AI in decision-making more broadly (e.g., [Foundation Models for Decision Making](#), [Generative Models for Decision Making](#)). Their scope within decision-making is extensive, encompassing diverse areas such as autonomous agents, robotics, and advancements in classical RL algorithms. In contrast, our workshop will concentrate on decision-making scenarios that are not merely optimization challenges but are fundamentally characterized by game-theoretic interactions and strategic behavior.

2. Tentative Schedule

The workshop will feature a dynamic full-day program designed to foster interaction and exchange between academia and industry. Our tentative schedule balances invited expert talks, contributed presentations, and interactive poster sessions, with flexibility to include a panel discussion pending speaker availability.

| Time | Session |
|---------------|---|
| 9:00 – 9:10 | Opening Remarks |
| 9:10 – 9:40 | Invited Talk 1 |
| 9:40 – 10:10 | Invited Talk 2 |
| 10:10 – 10:40 | <i>Coffee Break</i> |
| 10:40 – 11:40 | Oral Presentations (5 talks, 10 min presentation + 2 min Q&A each) |
| 11:40 – 13:00 | <i>Lunch Break</i> |
| 13:00 – 13:30 | Invited Talk 3 |
| 13:30 – 14:00 | Invited Talk 4 |
| 14:00 – 14:40 | <i>Coffee Break</i> |
| 14:40 – 15:10 | Invited Talk 5 |
| 15:10 – 15:50 | Poster Session (Interactive networking and discussion) |
| 15:50 – 16:00 | Best Paper Awards & Closing Remarks |

If the number of poster submissions is higher than anticipated, we will arrange two separate poster sessions. The additional session would take place in the morning for one hour, accompanied by a coffee break. The day's schedule would be adjusted by 30 minutes to accommodate this change.

3. Invited Speakers (Alphabetical Order)

The following five keynote speakers have already accepted our invitation.

- **Niklas Karlsson (Amazon):** Dr. Niklas Karlsson is a Senior Principal Research Scientist at Amazon Advertising, where he defines and directs the science vision for Amazon’s Demand Side Platform. Previously, he served as Chief Scientist and Vice President of R&D for Yahoo’s Demand Side Platform. He holds 50 issued patents in mobile autonomous robotics and online advertising and has received honors such as the Distinguished Alumni Award from UCSB and the Master Inventor Award from Yahoo. He is a Fellow of IEEE for his contributions to vSLAM and online advertising.
- **Tuomas Sandholm (CMU):** Tuomas Sandholm is Angel Jordan University Professor of Computer Science at Carnegie Mellon University, a serial entrepreneur, Co-Director of CMU AI, and Founder of the Electronic Marketplaces Laboratory. His research spans AI, economics, and operations research, with over 500 papers (h-index 97). He is a recipient of top honors like the IJCAI John McCarthy Award and a Fellow of ACM, AAAI, INFORMS, and AAAS.
- **Vijay Virkumar Vazirani (UC Irvine):** Vijay Virkumar Vazirani is a Distinguished Professor of Computer Science at the University of California, Irvine, and a leading figure in theoretical computer science. His research spans algorithms, complexity theory, and algorithmic game theory, with foundational contributions to approximation algorithms. He authored the seminal textbook *Approximation Algorithms*, which has defined the field for a generation of researchers. Professor Vazirani is a recipient of the Guggenheim Fellowship and is a Fellow of the ACM and AAAS for his fundamental contributions to algorithm design.
- **Zhenzhe Zheng (SJTU):** Zhenzhe Zheng is a Professor of Computer Science at Shanghai Jiao Tong University. After his PhD from SJTU and postdoctoral research at UIUC, his work at the intersection of economics, learning, and networking has resulted in over 90 publications. He has served as program chair for leading conferences such as KDD, IJCAI, and AAAI.
- **Song Zuo (Google):** Song Zuo is a Research Scientist at Google Research, specializing in Auction/Mechanism Design and Economics/Computation. He earned his PhD from IIIS, Tsinghua University, supported by a Google PhD Fellowship. His research has led to dozens of publications with over 1,000 citations, including a Best Paper Award at WWW 2024.

4. Organizers and Biographies

[Xiaotie Deng](#) (Peking University)

- Email: xiaotie@pku.edu.cn
- Bio: Xiaotie Deng is an ACM Fellow, IEEE Fellow, and a foreign member of Academia Europaea. He is a Chair Professor at Peking University and the Director of the Multi-Agent Intelligent Research Center at Peking University. His research interests include algorithmic game theory and Internet economics. He has published over 200 papers, with over 10,000 Google Scholar citations. He has served as chairman of multiple international academic conferences and initiated the globally circulating conference WINE.

[Jian Xu](#) (Alibaba Group)

- Email: xiyu.xj@alibaba-inc.com
- Bio: Jian Xu is the Chief Algorithm Architect of Alimama and a Senior Director of Alibaba Group. He leads the team in innovating and converting cutting-edge technologies into the Alimama ad platform, driving its continued success. He has published more than 80 papers in the areas of data mining, machine learning, and computational advertising. He also serves as

a program committee member and reviewer for various top academic conferences and journals in these areas.

[Fabrizio Silvestri](#) (University of Rome)

- Email: fsilvestri@diag.uniroma1.it
- Bio: Fabrizio Silvestri is a Full Professor and the coordinator of the Ph.D. in Data Science, at Dipartimento di Ingegneria informatica, automatica e gestionale (DIAG) of the University of Rome, La Sapienza. His research interests lie in Artificial Intelligence, and in particular, Fabrizio Silvestri deals with machine learning applied to web search problems and natural language processing. He is the author of more than 150 papers in international journals and conference proceedings. He is the holder of the “test-of-time” award at the ECIR 2018 conference for an article published in 2007. He is the holder of three best paper awards and other international awards. Fabrizio Silvestri spent eight years abroad in industrial research laboratories (Yahoo! and Facebook).

[Alireza Fallah](#) (Rice University)

- Email: afallah@rice.com
- Bio: Alireza Fallah is an Assistant Professor in the Department of Computer Science at Rice University and a member of the Ken Kennedy Institute. He earned his Ph.D. in Electrical Engineering and Computer Science from MIT in the summer of 2023, working with Asu Ozdaglar and Daron Acemoglu. Prior to joining Rice, he was a postdoctoral researcher at UC Berkeley, hosted by Michael Jordan. He also spent the fall of 2023 as the Gamelin Postdoctoral Fellow at the Simons Laufer Mathematical Sciences Institute (formerly MSRI), participating in the Mathematics and Computer Science of Market and Mechanism Design program. His research spans machine learning theory, economic theory, market and mechanism design, game theory, optimization, and privacy. His work integrates tools and insights from these fields to tackle challenges arising from the interplay between ML algorithms and human behavior.

[Yurong Chen](#) (INRIA Paris)

- Email: yurong.chen@inria.fr
- Bio: Yurong Chen is currently a postdoc at SIERRA-team, INRIA Paris, working with Prof. Michael I. Jordan. She earned her PhD in Computer Science at Peking University, where she was advised by Prof. Xiaotie Deng, and holds a bachelor’s degree in Applied Mathematics from the Hua Luogeng Honors Class at Beihang University. Her research focuses on the intersection of learning and game theory, especially on how strategic agents exploit information advantage against learning agents. She is a recipient of the Best Student Paper Award at WINE 2022.

[Brian Hu Zhang](#) (MIT)

- Email: zhangbh@csail.mit.edu
- Bio: Brian Hu Zhang is a postdoc at MIT CSAIL working with Profs. Constantinos Daskalakis and Gabriele Farina. He recently completed his PhD in Computer Science at Carnegie Mellon University under the supervision of Prof. Tuomas Sandholm. His current research interests lie in computational game theory, especially no-regret learning in games;

equilibrium computation in extensive-form games; solution concepts involving correlation, communication, and/or mediation; automated mechanism design; subgame solving; and connections to optimization.

Haoran Sun (Peking University)

- Email: sunhaoran0301@stu.pku.edu.cn
- Bio: Haoran Sun is a second-year Ph.D. student at Peking University, advised by Prof. Xiaotie Deng. His research interests include deep learning based mechanism design and the intersection of LLMs and game theory. He received his B.S. degree from Yuanpei College at Peking University in 2024.

Broad Experience in Academic Event Organization Our team members have extensive, hands-on experience organizing a wide spectrum of academic events, from major international conferences to specialized workshops, tutorials, and competitions. This diverse background provides us with deep expertise in managing all aspects of academic event planning and execution.

- **Conference Leadership:** Key roles include General Co-Chairs (ECIR 2025), General Chair (IJTCS 2020, WINE 2020), Conference Chair (ICALP 2020), Program Committee Chair (IEEE IPDPS 2020, SAGT 2018), and leadership on the steering committees for the [Conference on Web and Internet Economics \(WINE\)](#) and the [International Joint Conference on Theoretical Computer Science \(IJTCS\)](#).
- **Workshops and Tutorials:** Our experience includes serving as Hands-on Tutorial Chairs (KDD 2024) and organizing numerous workshops and tutorials, such as the [IR-RAG Workshop](#) (SIGIR 2025), the [Game AI Algorithms and Multi-Agent Learning \(GAAMAL\) Workshop](#) (IJCAI 2025), [The Economics of Data Tutorial](#) (EC 2025), and the [DLP-KDD workshop](#) (KDD 2019).
- **Competitions:** We have successfully organized major academic competitions, including the [Auto-bidding in Large-Scale Auctions Competition](#) (NeurIPS 2024) and the [Alibaba E-Commerce AI Challenge](#) (CIKM 2019).

This wealth of experience enables us to deliver a well-organized, engaging, and scientifically valuable event that fosters collaboration and advances the field.

5. Anticipated Audience Size

Based on ICLR’s broad and diverse community spanning academia and industry, as well as historical attendance patterns from previous ICLR workshops, we anticipate approximately 40 submissions to our workshop. We further project an in-person and virtual audience of around 100 attendees, including authors, researchers, practitioners, and interested participants from related fields.

6. Advertising and Outreach Plan

Our advertising and outreach plan is designed to maximize visibility and engagement across both academic and industry communities. Leveraging our organizing team’s global presence across three continents and six institutions, we have a uniquely broad network for promotion.

Academic Outreach: We will utilize the academic leadership of **Peking University, MIT, INRIA, Rice University, and the University of Rome La Sapienza** to disseminate the Call for Papers (CFP) through departmental and laboratory **mailing lists, PhD forums**, and collaborations with **affiliated institutions in Asia, North America, and Europe**. Our Program Committee members will further amplify the call within their own extensive research circles.

Industry Engagement: Our partners, **Alibaba Cloud** and **DAMO Academy**, are well-known brands in the AI community. We will leverage their networks to promote the workshop through **official newsletters**, targeted **social media channels**, and outreach to their **industry partners**.

Incentives and Support: To encourage high-quality submissions and broad attendance, we will leverage a dedicated outreach budget provided by Alibaba Group. This will allow us to offer **awards and internship invitations** for top papers and **souvenirs** for all participants.

Digital Presence & Timeline: Preparations are already underway on the workshop website and promotional poster. The website will be optimized for both desktop and mobile users, and the poster will be shared widely. To broaden our reach, we will conduct cross-platform promotion on **X (formerly Twitter)**, **LinkedIn**, and specialized platforms like **WeChat** and **Red** for Chinese-speaking communities. We will launch the official website and release the CFP within one week of acceptance, followed by regular updates and reminders until the submission deadline.

We are confident these combined efforts will ensure a successful workshop with high visibility and strong engagement.

7. Diversity Commitment

We are deeply committed to fostering a diverse, inclusive, and globally representative workshop. This commitment is a cornerstone of our planning, reflected in the composition of our organizing committee and our selection of invited speakers. We believe that bringing together a wide variety of perspectives is essential for driving innovation and enriching the intellectual discourse for all participants.

Diversity in Organization Our organizing committee has been intentionally assembled to represent a broad spectrum of geographical locations, professional perspectives, and career stages.

The committee spans **three continents**, including Xiaotie Deng, Jian Xu, and Haoran Sun from Asia; Fabrizio Silvestri and Yurong Chen from Europe; and Alireza Fallah and Brian Hu Zhang from North America. The seven organizers are affiliated with six different institutions, including top-tier academic institutions (PKU, University of Rome La Sapienza, INRIA, Rice University, MIT) and a leading industrial research lab (Alibaba Group).

This group also embodies a **strong bridge between academia and industry**. Prof. Deng, Prof. Fallah, Dr. Chen, and Dr. Zhang represent the academic frontier, while Jian Xu provides senior leadership experience from an industrial perspective. Prof. Silvestri further strengthens this connection, having held research scientist positions at Facebook and Yahoo before his current role as a full professor.

Furthermore, our committee includes voices from **every major academic career stage**, from senior professors (Xiaotie Deng and Fabrizio Silvestri) to an assistant professor (Alireza Fallah), two postdoctoral researchers (Yurong Chen and Brian Hu Zhang), and a PhD student (Haoran

Sun). With senior researchers providing overall guidance and attracting leading speakers, and junior members handling coordination and logistics, this diversity in experience ensures that our workshop is inclusive, relevant, and guided by a comprehensive, forward-looking vision.

Diversity in Speakers We are also committed to inviting world-leading experts from both academia and industry. Our five confirmed speakers represent diverse career trajectories and geographical regions.

From academia, we are honored to host highly distinguished senior scholars **Tuomas Sandholm** (CMU) and **Vijay Virkumar Vazirani** (UC Irvine) from the United States, alongside **Zhenzhe Zheng** (Shanghai Jiao Tong University), an accomplished professor who provides a key perspective from Asia.

Complementing these academic viewpoints are two outstanding speakers from industry. **Niklas Karlsson** (Amazon), a Senior Principal Research Scientist and IEEE Fellow, offers a wealth of experience from his senior leadership roles in top technology companies. Representing the next generation of industrial research leaders, **Song Zuo** (Google Research), a recent Google PhD Fellow and Best Paper Award winner, will provide fresh insights into cutting-edge research and its application. This lineup ensures a balanced dialogue between foundational academic principles and their high-stakes application in industry.

Through this deliberate composition of our organizers and speakers, we are confident the workshop will provide a vibrant, globally relevant, and intellectually inclusive environment for all attendees.

8. Virtual Access

Recognizing the global and diverse nature of the ICLR community and acknowledging the challenges posed by travel constraints, time zone differences, and unpredictable attendance, we are committed to ensuring that the intellectual content and discussions of our workshop remain accessible to all interested participants, regardless of their ability to attend in person.

All invited talks and oral presentations will be **professionally recorded** (with speaker consent) and **made publicly available** on our official workshop website within one week of the event. For poster presenters, we will invite submission of 3-minute “elevator pitch” videos alongside their PDF posters, which will be hosted alongside presentation recordings to facilitate asynchronous engagement. Additionally, we will **maintain a curated, long-term workshop webpage** featuring: (1) speaker slides, (2) accepted paper abstracts and links to non-archival PDFs, (3) discussion summaries or key takeaways from Q&A and poster sessions, and (4) a public Slack/Discord channel (or forum) for ongoing conversation, resource sharing, and community building beyond the workshop day.

To further extend impact, we are exploring the possibility of curating a post-workshop blog series or partnering with an open-access venue to publish selected reflections, case studies, or “lessons learned” contributions — particularly those from industry practitioners. This hybrid-access model, powered by Alibaba Group’s technical infrastructure and Peking University’s academic stewardship, ensures that our workshop’s value is not confined to a single day or physical location, but becomes a lasting, globally accessible resource for the broader ML community. Our technical requirements include standard audio-visual support such as presentation projection and high-speed internet access.

9. Program Committee Members

Our review process is designed to provide authors with timely and high-quality feedback. Each submission will receive at least three independent reviews, with a maximum load of five papers per reviewer. The Program Committee (PC) comprises experts in Mechanism Design, Computational Economics, Algorithmic Game Theory, Machine Learning, and Large Language Models, selected for their significant contributions and expertise. We will strictly adhere to the ICLR conflict-of-interest policy. To manage review load and maintain our timeline, we will confirm PC availability upon launching the submission platform, recruit emergency reviewers, and invite submitting authors to contribute to the review process.

Below is a confirmed list of PC members (total: 26). We are actively expanding this list to ensure broad expertise and diverse perspectives.

- Dirk Bergemann (Yale University)
- Zhiyi Huang (HKU)
- David C. Parkes (Harvard University)
- Alex Smolin (TSE)
- Yiding Feng (HKUST)
- Hanrui Zhang (CUHK)
- Ningyuan Li (PKU)
- Dagui Chen (Alibaba)
- Xuanzhi Xia (Tsinghua University)
- Yutong Yin (Northwestern University)
- Weiran Shen (RUC)
- Jason Hartline (Northwestern University)
- Fan Yao (UNC)
- Yuwei Cheng (University of Chicago)
- Yixin Tao (SUFU)
- Yifeng Teng (Google)
- Dengji Zhao (ShanghaiTech)
- Jason Hartline (Northwestern University)
- Vijay V. Vazirani (UC Irvine)
- Zhilin Zhang (Alibaba)
- Zhijian Duan (Independent Researcher)
- Chuan Yu (Alibaba)
- Zihe Wang (RUC)
- Jibang Wu (University of Chicago)
- Yuwei Cheng (University of Chicago)
- Nikhil R Devanur (Amazon)

References

- [1] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [2] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1-2):166–196, 2001.
- [3] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [4] Yadati Narahari. *Game theory and mechanism design*, volume 4. World Scientific, 2014.
- [5] Kathleen M Eisenhardt and Mark J Zbaracki. Strategic decision making. *Strategic management journal*, 13(S2):17–37, 1992.
- [6] Charles R Schwenk. Strategic decision making. *Journal of management*, 21(3):471–493, 1995.
- [7] Said Elbanna. Strategic decision-making: Process perspectives. *international Journal of Management reviews*, 8(1):1–20, 2006.

[8] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[9] Jérémie Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

[10] Michael P. Wellman. Methods for empirical game-theoretic analysis. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI'06)*, pages 1553–1557. AAAI Press, 2006.

[11] Martin L. Weitzman. Prices vs. quantities. *The Review of Economic Studies*, 41(4):477–491, 1974.

[12] Anjan V Thakor. The design of financial systems: An overview. *Journal of Banking & Finance*, 20(5):917–948, 1996.

[13] Qingpeng Cai, Aris Filos-Ratsikas, Pingzhong Tang, and Yiwei Zhang. Reinforcement mechanism design for e-commerce. In *Proceedings of the 2018 World Wide Web Conference*, pages 1339–1348, 2018.

[14] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3354–3364, 2021.

[15] Esteban A Soto, Lisa B Bosman, Ebisa Wollega, and Walter D Leon-Salas. Peer-to-peer energy trading: A review of the literature. *Applied energy*, 283:116268, 2021.

[16] Sangdon Park, Joohyung Lee, Sohee Bae, Ganguk Hwang, and Jun Kyun Choi. Contribution-based energy-trading mechanism in microgrids for future smart grid: A game theoretic approach. *IEEE Transactions on Industrial Electronics*, 63(7):4255–4265, 2016.

[17] Mohsen Khorasany, Reza Razzaghi, and Amin Shokri Gazafroudi. Two-stage mechanism design for energy trading of strategic agents in energy communities. *Applied Energy*, 295:117036, 2021.

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[19] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[20] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[21] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

[22] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[23] Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming*, pages 19–36. Springer, 2003.

[24] Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multi-agent reinforcement learning. In *International conference on machine learning*, pages 5571–5580. PMLR, 2018.

[25] Junqi Jin, Chengru Song, Han Li, Kun Gai, Jun Wang, and Weinan Zhang. Real-time bidding with multi-agent reinforcement learning in display advertising. In *Proceedings of the 27th ACM international conference on information and knowledge management*, pages 2193–2201, 2018.

[26] Ariyan Bighashdel, Yongzhao Wang, Stephen McAleer, Rahul Savani, and Frans A Oliehoek. Policy space response oracles: A survey. *arXiv preprint arXiv:2403.02227*, 2024.

[27] Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning: Advances in differentiable economics. *Journal of the ACM*, 71(1):1–53, 2024.

[28] Michael Curry, Tuomas Sandholm, and John Dickerson. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2633–2641, 2023.

[29] Zhijian Duan, Haoran Sun, Yurong Chen, and Xiaotie Deng. A scalable neural network for dsic affine maximizer auction design. *Advances in Neural Information Processing Systems*, 36:56169–56185, 2023.

[30] Neehar Peri, Michael Curry, Samuel Dooley, and John Dickerson. Preferencenet: Encoding human preferences in auction design with deep learning. *Advances in Neural Information Processing Systems*, 34:17532–17542, 2021.

[31] Jad Rahme, Samy Jelassi, and S. Matthew Weinberg. Auction learning as a two-player game. In *International Conference on Learning Representations*, 2021.

[32] Jad Rahme, Samy Jelassi, Joan Bruna, and S Matthew Weinberg. A permutation-equivariant neural network architecture for auction design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5664–5672, 2021.

[33] Mai Pham, Vikrant Vaze, and Peter Chin. Advancing differentiable economics: A neural network framework for revenue-maximizing combinatorial auction mechanisms. *arXiv preprint arXiv:2501.19219*, 2025.

[34] Zhijian Duan, Jingwu Tang, Yutong Yin, Zhe Feng, Xiang Yan, Manzil Zaheer, and Xiaotie Deng. A context-integrated transformer-based neural network for auction design. In *International Conference on Machine Learning*, pages 5609–5626. PMLR, 2022.

[35] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.

[36] Haoran Sun, Yusen Wu, Yukun Cheng, and Xu Chu. Game theory meets large language models: A systematic survey. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2025.

[37] Ruitao Zhu, Yangsu Liu, Dagui Chen, Zhenjia Ma, Chufeng Shi, Zhenzhe Zheng, Jie Zhang, Jian Xu, Bo Zheng, and Fan Wu. Contextual generative auction with permutation-level externalities for online advertising. *arXiv preprint arXiv:2412.11544*, 2024.

[38] Kefan Su, Yusen Huo, Zhilin Zhang, Shuai Dou, Chuan Yu, Jian Xu, Zongqing Lu, and Bo Zheng. Auctionnet: A novel benchmark for decision-making in large-scale games. *Advances in Neural Information Processing Systems*, 37:94428–94452, 2024.

[39] Jie Sun, Tianyu Zhang, Houcheng Jiang, Kexin Huang, Chi Luo, Junkang Wu, Jiancan Wu, An Zhang, and Xiang Wang. Large language models empower personalized valuation in auction. *arXiv preprint arXiv:2410.15817*, 2024.

[40] Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *Nature Human Behaviour*, pages 1–11, 2025.

[41] Nunzio Lorè and Babak Heydari. Strategic behavior of large language models and the role of game structure versus contextual framing. *Scientific Reports*, 14(1):18490, 2024.

[42] Caoyun Fan, Jindou Chen, Yaohui Jin, and Hao He. Can large language models serve as rational players in game theory? a systematic analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17960–17967, 2024.

[43] Xiachong Feng, Longxu Dou, Ella Li, Qinghao Wang, Haochuan Wang, Yu Guo, Chang Ma, and Lingpeng Kong. A survey on large language model-based social agents in game-theoretic scenarios. *arXiv preprint arXiv:2412.03920*, 2024.

[44] Soheil Feizi, MohammadTaghi Hajiaghayi, Keivan Rezaei, and Suho Shin. Online advertisements with llms: Opportunities and challenges. *arXiv preprint arXiv:2311.07601*, 2023.

[45] Paul Duetting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. Mechanism design for large language models. In *Proceedings of the ACM Web Conference 2024*, pages 144–155, 2024.

[46] MohammadTaghi Hajiaghayi, Sébastien Lahaie, Keivan Rezaei, and Suho Shin. Ad auctions for llms via retrieval augmented generation. *arXiv preprint arXiv:2406.09459*, 2024.

[47] Avinava Dubey, Zhe Feng, Rahul Kidambi, Aranyak Mehta, and Di Wang. Auctions with llm summaries. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 713–722, 2024.

[48] Ermis Soumalias, Michael J Curry, and Sven Seuken. Truthful aggregation of llms with an application to online advertising. *arXiv preprint arXiv:2405.05905*, 2024.