# Project Instructions
# CMSC 35401: The Interplay of Economics and Machine Learning
# (Winter'24)

## General Instructions

The project will count for 45% of your total grade.

The goal of this project is to apply the basic concepts and theory you learned from this course to a research project of your own interest and develop some nonvel and non-trivial solutions. A successful project could be, e.g., using the concepts or theory you learned to formulate a research question in your own research field and then try to provide an (even partial) solution. **Please notice that though the course is theory-oriented, your course project does not need to be — i.e., applied project is equally welcome!**

You are encouraged to form a team of 2-4 members to complete the project together. If needed, we are happy to help you with finding team members — please feel free to email me or the TA if so (we can also create a discord or slack discussion channel if there is a lot of demand). If you really want to do a project alone, please come to talk with us first since we would like to make sure the project is doable for one person and also has sufficient content — remember that it counts for a big portion of of your grade.

Please feel free to discuss with me or the TA Alec regarding, e.g., identifying a topic, references, ideas for the project, though these are not mandatory.

**Notice**: we understand that research is risk and unpredictable. So your project does not have to show prefect results — in fact, you can have a successful project by showing all the thing you have tried and *why they failed*. The primary goal is for you to exercise the concepts and techniques you learned in an interesting research domain. As long as we see that you have tried out non-trivial methods (which is usually easy to see), there should not be much to worry.

## Timeline and Presentation Formats

**Project Proposal** (5 **points**) — this is primarly to make sure that you do have a team and have a problem to work on. Submission is an *one-page* file that describes what are your team members, what you would like to work on, and some initial thoughts/survey. It is due by Jan 27'th 9 pm. You will submit it to Gradescope — notably, every team member needs to submit it, even though you will be submitting the same document (since this makes it easier for giving grades on Gradescope).

**Project Presentation** (15 **points**) — each project will be given about $10 \sim 15$ mins to be presented during the last class of this course (i.e., 2/29 Thursday).

**Project Report (25 points)** — majority of the score is assigned to the final report in PDF format, which should be at most 7 pages with single-column, 11 pt fonts and 1 inch margin. This is due by March 7'th (03/07) 9 pm to Gradescope as well — notably, every team member needs to submit it, even though you will be submitting the same report.

## Suggested Project Topics (biased due to the instructor's research experience)

Importantly, we would like to reiterate that we strongly encourage you to identify your own topics for the project. The suggested topics below is supposed to serve as your second resort, if you were not able to find your own project topics. Moreover, please feel free to discuss with the instructor for suggestions/comments on your ideas and for additional sources of references. Some interesting projects from this course have led to published papers at top ML conferences such as NeurIPS.

1. Understanding the multi-agent interactions among Large Language Models (LLMs) or understand LLM's strategic reasoning capabilities.

   There has been much recent studies on this, but there is much room to improve. For instance, a recent ICLR 2024 paper [DVJ+24] propose to use negotiation games as benchmarks for evaluating strategic behaviors of LLMs. This recent UIST best paper [POC+23] tries to understand what kind of society it will be if every agent in the society is an LLM.

2. Designing/developing AI agents for playing games that require *strategic reasoning* + *natural language*. One such example game is the *Diplomacy* game. A recent Science paper by Meta developed a human-level AI agent for playing Diplomacy that combine game-theoretic optimization with natural language [FBB+22]. They did not even use LLM yet and are using just basic NLP generation technique; another recent paper solves werewolf (similar to Mafia) using similar high-level approaches but utilizing LLMs. There are still much room to improve upon in this space.

   Besides the entertainment games mentioned above, I think a even much more impactful direction is to design such NLP-powered strategic AI agent for many real-world applications, such as negotiation and bargaining.

3. Understanding the effort tension between foundation model training and fine-tuning, which are often down by different parties with their different goals. See a recent study in [LKH23] on modeling and understanding such strategic interactions.

4. Rethinking the online content creation ecosystem in the era of generative AI.

   The advances of generative AIs are great, but also poses significant threats to human content creators. Due to the low cost of creating contents by generative AIs as well as their creativity on certain aspects (despite limitations otherwise), the competition of content creation may incentivize creators to overwhelmingly only use AI to generate contents. This, however, is disastrous to the generative AI technology itself as its success relies on training with large amount of authentic data. In fact, it is not difficult to imagine that the success of future technology firms will largely depend on the richness of the authentic data they possess for updating their systems or generative models. Concerns above raise novel research questions such as how to foster healthy competition between AI- generated and human-generated contents, and how to incentivize desirable content generation from a platform's perspective. This project will start a systematic investigation of these research questions.

5. The valuation, pricing and acquisition of Data.

On one hand, as foundation models become increasingly capable of generating synthesized images, dialogues and videos, they will significantly threaten original content creators' incentives for generating any new data since more and more Internet traffic will divert from human created contents to AI-generated contents. On the other hand, sustained human creation of high-quality authentic content data are vital to the development of powerful foundation models. This tension is a key reason that we have seen recent lawsuits of New York Times against Microsoft [Nyt23], Getty Images against Stability AI [Get23], in order to seek a proper share of commercial success of the model. Meanwhile, OpenAI has contracted with news vendors [ap23] and image curators [shu23] to gain access to their content. Hence, it becomes an important incentive design problem to determine the optimal acquisition mechanism of content data and to share the platform's revenue among data vendors in order to ensure the sustainable development of foundation models. Research questions along this line are extremely rich and interdisciplinary, and necessitie deep understandings on the valuation of data in machine learning (e.g., data Shapley [GZ19]), data privacy control, system-level efficiency of data discovery, as well as various economic paradigms such as contract/mechanism/information design

6. Incentive design in Bitcoin mining. One recent hot trend is to understand gaming phenomenon in bitcoin mining and how to design effective mechanisms to incentive the right mining behaviors (e.g., transaction fee design). See, e.g., [CS21, Rou21]. You can do a survey or study new questions on this frontier.

7. Use online learning algorithms (e.g., multiplicative weight) to solve structured games such as Stackelberg games, contract design problems, principal-agent problems and security games.

   See this paper [GHWX22] for a general framework, which includes Stackelberg games, contract design, Bayesian persuasion as special cases. Another interesting class of games is the high impactful *security game*, which is a fundamental resource allocation game with significant real-world impact. See, e.g., this paper [Xu16] for an introductory reading. Previous research mostly used linear programming or OR techniques to solve security games. You may think about using online learning algorithm to solve the game, and see whether it is faster than classical OR-based algorithms. Multiplicatve weight is a very powerful idea for online leanring. This survey paper [AHK12] shows how powerful multiplicative weight is.

8. The study of convergence of no-regret learning algorithms to equilibriua in structured games

   We looked at convergence of no-regret learning algorithm in zero-sum games, and some of you figured out that it also provably converges in $2 \times 2$ matrix games. Is there any other structured settings — e.g., a game arised in your research domain or a special game classes such as Stackelberg games, principal-agent problems [GHWX22], congestion games, or dominance elimination solvable games, etc. — that no-regret learning algorithm also provably or empirically converges to equilibrium? If the algorithm does not convergence, what is the patten of the its trajectory?

   There has been extensive literature on this frontier for zero-sum games [PP16, MJS19], congestion games [CXFD22], dominance-elimination solvable games [WXY22, WKBJ22], monotone games [COZ22]. However, many problems remain open. For example, the regret bound of [CXFD22] leaves significant room for improvement. An interesting open problem from [WXY22] is whether there exists a no-regret learning algorithm that provably converges to rationalizable equilibrium in two-player games.

9. You may ask a more applied research question. That is, can you design fully applied algorithm — e.g., a deep learning algorithm by carefully constructing the network architecture to suit game-theoretic applications — to solve the listed games above or beyond.

## References


[AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of computing*, 8(1):121–164, 2012.

[ap23] Chatgpt-maker openai signs deal with ap to license news stories. Online, 2023.

[COZ22] Yang Cai, Argyris Oikonomou, and Weiqiang Zheng. Finite-time last-iterate convergence for learning in multi-player games. *Advances in Neural Information Processing Systems*, 35:33904–33919, 2022.

[CS21] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. *arXiv preprint arXiv:2111.03151*, 2021.

[CXFD22] Qiwen Cui, Zhihan Xiong, Maryam Fazel, and Simon S Du. Learning in congestion games with bandit feedback. *arXiv preprint arXiv:2206.01880*, 2022.

[DVJ+24] Tim R Davidson, Veniamin Veselovsky, Martin Josifoski, Maxime Peyrard, Antoine Bosselut, Michal Kosinski, and Robert West. Evaluating language model agency through negotiations. *arXiv preprint arXiv:2401.04536*, 2024.

[FBB+22] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

[Get23] Getty images (us), inc. v. stability ai, inc. (1:23-cv-00135). Online, 2023.

[GHWX22] Jiarui Gan, Minbiao Han, Jibang Wu, and Haifeng Xu. Optimal coordination in generalized principal-agent problems: A revisit and extensions. *arXiv preprint arXiv:2209.01146*, 2022.

[GZ19] Amirata Ghorbani and James Zou. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*, pages 2242–2251. PMLR, 2019.

[LKH23] Benjamin Laufer, Jon Kleinberg, and Hoda Heidari. Fine-tuning games: Bargaining and adaptation for general-purpose models. *arXiv preprint arXiv:2308.04399*, 2023.

[MJS19] Eric V Mazumdar, Michael I Jordan, and S Shankar Sastry. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *arXiv preprint arXiv:1901.00838*, 2019.

[Nyt23] The new york times company v. microsoft corporation (1:23-cv-11195). Online, 2023.

[POC+23] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22, 2023.

[PP16] Christos Papadimitriou and Georgios Piliouras. From nash equilibria to chain recurrent sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 227–235, 2016.

[Rou21] Tim Roughgarden. Transaction fee mechanism design. *ACM SIGecom Exchanges*, 19(1):52–55, 2021.

[shu23] Shutterstock partners with openai and leads the way to bring ai-generated content to all. Online, 2023.

[WKBJ22] Yuanhao Wang, Dingwen Kong, Yu Bai, and Chi Jin. Learning rationalizable equilibria in multiplayer games. 2022.

[WXY22] Jibang Wu, Haifeng Xu, and Fan Yao. Multi-agent learning for iterative dominance elimination: Formal barriers and new algorithms. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 543–543. PMLR, 02–05 Jul 2022.

[Xu16] Haifeng Xu. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 497–514, 2016.