Summary of Past Research

In the rapidly evolving field of machine learning, my research endeavors to pioneer efficient and robust algorithms at the forefront of technological advancement with solid theoretical guarantees.

- **Robust Reinforcement Learning in Dynamic Environments**.
  - *Environment shifting in Markov games [6].* In practical scenarios, large-scale models are often trained in simulated environments before immigrating to real-world settings. However, this strategy often experiences a considerable performance reduction due to the environment shifting, especially in the stochastic games. To handle this issue, I proposed a theory to characterize the equilibrium considering the worst-case scenario [6] with a corresponding algorithm that can efficiently learn the optimal strategy in such a highly uncertain environment.
  - *Training in dynamic environments [9, 4].* When training a practical model, ensuring the data always follows a certain distribution is usually hard. My efforts have been directed at theoretically explaining how the performance of SGD would be influenced in such a practical environment. More interestingly, we found applying large batch sizes is always helpful, especially when the data is highly correlated [4].
- **Design of Efficient Algorithms**.
  - *Efficient adversarial algorithms [2].* Optimization algorithms are crucial in training large-scale machine learning models. However, the training procedure usually has more complicated structures in many scenarios, such as training adversarial or attacking-robust models. To circumvent the limitation of traditional algorithms in these situations, I proposed an accelerating training strategy for the adversarial training pipeline [2]. Especially in the modern cybersecurity field, adversarial attacking techniques easily fool the large vision model. Our approach allows us to rapidly train a robust model for defending such an attacking strategy.
  - *Efficient reinforcement learning algorithms [10, 5, 1, 3].* Reinforcement learning (RL) is a pivotal tool in developing highly autonomous agents. Despite their efficacy, training an RL model requires millions of interactions with the environment, which is typically expensive and unfeasible. It motivates me to design algorithms significantly reducing the interactions needed for training reinforcement learning models. Notable works include evaluating the performance of an RL model [10], learning an equilibrium agent in a competitive environment [1, 3], and training an RL model to achieve the optimal performance [5].
- **Trustworthy Scientific Models**.
  - *Training physics-informed models [8, 7].* The fluid flow prediction task requires the model to output robust and high-accurate field prediction on unseen or out-of-distribution parameters. So, embedding the physical simulation result in the neural network model is preferred. However, most of the physics simulator doesn't support auto-differentiation. While collaborating with Lawrence Livermore National Security, I addressed this intricate challenge of training such a hybrid scientific model using zeroth-order gradient estimation and obtained comparable performance.

As I advance in my career, I am committed to further exploring these fascinating intersections, contributing to the academic community, and practical implementations in various domains.

Publications

[1] Ziyi Chen, Shaocong Ma, and Yi Zhou. Sample efficient stochastic policy extragradient algorithm for zero-sum markov game. In *International Conference on Learning Representations*, 2021.

[2] Ziyi Chen, Shaocong Ma, and Yi Zhou. Accelerated proximal alternating gradient-descent-ascent for nonconvex minimax machine learning. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 672–677. IEEE, 2022.

[3] Ziyi Chen, Shaocong Ma, and Yi Zhou. Finding correlated equilibrium of constrained markov game: A primal-dual approach. *Advances in Neural Information Processing Systems*, 35:25560–25572, 2022.

[4] Shaocong Ma, Ziyi Chen, Yi Zhou, Kaiyi Ji, and Yingbin Liang. Data sampling affects the complexity of online sgd over dependent data. In *Uncertainty in Artificial Intelligence*, pages 1296–1305. PMLR, 2022.

[5] Shaocong Ma, Ziyi Chen, Yi Zhou, and Shaofeng Zou. Greedy-gq with variance reduction: Finite-time analysis and improved complexity. In *International Conference on Learning Representations*, 2020.

[6] Shaocong Ma, Ziyi Chen, Shaofeng Zou, and Yi Zhou. Decentralized robust v-learning for solving markov games with model uncertainty. *The Journal of Machine Learning Research (JMLR)*, 2023.

[7] Shaocong Ma, James Diffenderfer, Bhavya Kailkhura, and Yi Zhou. End-to-end mesh optimization of a hybrid deep learning black-box pde solver. *NeurIPS 2023 (ML4PS Workshop)*, 2023.

[8] Shaocong Ma, James Diffenderfer, Bhavya Kailkhura, and Yi Zhou. When non-differentiable pde solver meets deep learning: Partially differentiable learning for efficient fluid flow prediction. *(Under review)*, 2023.

[9] Shaocong Ma and Yi Zhou. Understanding the impact of model incoherence on convergence of incremental SGD with random reshuffle. In *International Conference on Machine Learning*, pages 6565–6574. PMLR, 2020.

[10] Shaocong Ma, Yi Zhou, and Shaofeng Zou. Variance-reduced off-policy tdc learning: Non-asymptotic convergence analysis. *Advances in neural information processing systems*, 33:14796–14806, 2020.