Research Statement: Statistical Game Theory

Arun Sai Suggala (asuggala@cs.cmu.edu)

Game theory and statistics are two huge scientific disciplines that have played a significant role in the development of a wide variety of fields, including computer science, natural sciences, and social sciences. Traditionally, game theory has been used for decision making in strategic environments. For example, in economics, it is often used for designing auctions and for decision making in competitive markets. In computer science, it has found applications in numerous sub-fields such as distributed computing, network security, robotics, multi-agent systems, where multiple self-interested parties interact with each other. Unlike game theory, statistics has traditionally been used for reasoning in non-strategic and non-adversarial environments. In particular, statistics is concerned with the analysis and interpretation of data generated by some stationary non-reactive source. For example, in numerous fields such as astronomy, biostatistics, business analytics, epidemology, finance, statistical analysis and estimation is often performed on data generated from non-reactive sources. Due to the contrasting settings in which game theory and statistics are often studied, these two disciplines have traditionally been regarded as disparate research areas. However, there is a great degree of commonality between the two fields. A surprising range of developments in classical and modern statistics have a game theoretic component to them:

- **Classical Developments.** Classically, the mathematical philosophy of statistics, in particular frequentist statistics, was concerned about strategic considerations. It posits that the source of samples seen by the statistician is potentially adversarial. This resulted in the rich theory of minimax statistical estimation and games. In these games, statistical estimation problems are framed as two-player games in which nature adversarially selects a distribution that makes it difficult for a statistician to perform the estimation.
- **Modern Developments.** Boosting algorithms, which are often regarded as best off-the-shelf classifiers, can be viewed as playing a zero-sum game against a weak learner. To allow for various departures of "test environment" from "train environments", the emerging field of robust machine learning allows for adversarial manipulation of the train or test environments. An emerging class of density estimators in modern machine learning use an adversarial "critic" of the density estimator to improve the final density estimation.

The common theme among these classical and modern developments is an interplay between statistical estimation and two-player games.

I aim to bring together statistical machine learning and game theory, and study the interplay between the two fields. In particular, I'm interested in studying machine learning and statistical problems from a game theoretic perspective. Despite the many commonalities between the two fields, the game theoretic perspective of many machine learning and statistical problems is often ignored due to various analytical and computational reasons. One of the unpleasant facts about many games arising in machine learning and statistics is that they are generally much too big and too difficult to solve than those typically arising in economics and computer science. For example, consider the problem of minimax statistical estimation, which can be viewed as a game between statistician and nature. The action space of the statistician in this game is the set of all functions, which is an infinite dimensional space. Existing algorithmic tools from game theory are inefficient for solving this game. Consequently, statisticians have often ignored the game theoretic viewpoint while designing minimax estimators.

Setting aside these analytical and computational caveats, the game theoretic perspective provides tremendous value, and comes with several benefits. It can help us reason about and construct optimal solutions for the wide range of statistical problems described above. As an example, consider again the problem of minimax statistical estimation. Existing approaches for designing minimax estimators often rely on prior knowledge and require a deeper understanding of the problem at hand. This process is very time consuming, and often requires decades of research on the problem; for example, designing the popular LASSO estimator required decades of research on sparse estimation. In contrast, the game theoretic perspective can help us come up with algorithmic approaches which can automate the process of designing minimax estimators. Such algorithmic approaches can be of tremendous value to the statisticians, as they can aid them in building minimax estimators. As another example, consider robust machine learning. Existing approaches for constructing robust models often rely on heuristics and are not guaranteed to return an optimal solution. In contrast, the game theoretic viewpoint of robust machine learning provides us a wide array of tools for constructing robust models which can withstand adversarial manipulations better than existing approaches.

I'm interested in developing new algorithmic and analytic tools in **game theory** which address the caveats mentioned above, and help us study **machine learning** and **statistical** problems from a game theoretic viewpoint. During my PhD, I have primarily focused on minimax statistical estimation and boosting. My work on minimax statistical estimation has provided efficient techniques for algorithmically building minimax estimators [1, 2, 3]. For various fundamental problems such as regression and entropy estimation, our algorithmic minimax estimators match, if not beat, the performance of existing minimax estimators designed by statisticians. My work on boosting has aimed to improve its performance and bring it closer to the performance of neural networks [4].

1 Minimax Statistical Estimation

For decades, minimax statistical estimation has been crucial for the development of frequentist statistics, as it aids statisticians in picking estimators that work well even under the worst circumstances. Traditional approaches for solving minimax statistical games often require a deeper understanding of the problem, can be time-consuming, and do not provide concrete guidelines for designing such estimators for new estimation problems [5]. So algorithmic approaches that automate this process can be of immense help to statisticians. In this section, I will describe my work on automating the process of designing minimax estimators. At a high-level, I approach this problem from a game theoretic perspective and aim to develop efficient algorithms for finding a Nash equilibrium of the statistical game. However, this approach poses a number of challenges:

- The statistical game is often not convex-concave. Consequently, algorithms developed for convex-concave games are not guaranteed to find an equilibrium of the statistical game.
- In contrast to the typical zero-sum games studied in economics and computer science, the set of all possible moves of the statistician is extremely large.

Tackling these issues requires advances in a number of sub-fields including **online nonconvex learning**, **derivative-free optimization** and **Bayesian estimation**. My previous and ongoing work in PhD has primarily focused on tackling these issues.

Online Nonconvex Learning (previous work). Over the years, online learning has played a key role in advancing game theory and machine learning. It has provided some of the most successful optimization methods used in game theory. However, most of the existing works on online learning have focused on the convex setting. In contrast, solving minimax statistical games requires efficient and optimal algorithms for online learning in the more general nonconvex setting. In [2, 3], we develop algorithms that satisfy these desiderata. Our algorithms reduce the problem of online learning to the well-studied problem of offline optimization. Since the latter problem is well understood even in the nonconvex setting, our algorithms can be efficiently implemented for many problems of interest by leveraging the rich body of work on global optimization. We note that our algorithms are of independent interest and have several applications beyond statistical games considered here. The most important of these are their applications to online learning in bandit setting, contextual bandits, and solving nonconvex-nonconcave games arising in other areas such as robust machine learning, adversarial density estimation.

Statistical Games (previous work). As previously mentioned, a critical distinction of the statistical games, in contrast to other games, is that the action space of the statistician is extremely large. In [1], we show that it is nonetheless possible to finesse this technical caveat and solve the statistical game using the online learning algorithms developed in [2, 3], provided we are given access to two optimization subroutines: a Bayes estimator subroutine which outputs a Bayes estimator corresponding to any given prior, and a subroutine which computes the worst-case risk of any given estimator. Given access to these two subroutines, we show that our algorithm outputs a minimax estimator. In Table 1, we present the performance of our algorithmically constructed estimator for the problem of entropy estimation. This problem lies at the core of many applications of information theory to data analysis. It can be seen that our estimator beats the classical plugin estimator and has similar performance as the best known estimator recently developed by [6]. More empirical results on other problems such as linear regression, covariance estimation can be found in [1].

Table 1: Worst-case risk of various estimators for entropy estimation. Here n is the number of samples and d is the dimension.

	Worst-case Risk			
	d = 40		**d = 80**	
Estimator	**n = 20**	**n = 40**	**n = 40**	**n = 80**
Plugin estimator (Classical)	0.8909	0.2710	0.9142	0.2899
JVHW [6]	0.2699	0.0648	**0.1755**	0.0974
Algorithmic Estimator [1]	**0.2320**	0.0822	0.2539	0.0672

Derivative free optimization (future Work). Scaling the above algorithm to high dimensional problems requires efficient implementation of the two optimization subroutines described above. Implementing the Bayes estimator subroutine is easy, as one can rely on the vast array of tools developed by Bayesians for efficiently computing Bayes estimators. In contrast, implementing the subroutine for computing the worst-case risk can be computationally expensive, as it requires the minimization of a nonconvex objective. To efficiently implement this subroutine in high dimensions, we need a derivative-free optimization (DFO) technique which satisfies the following desiderata: (a) scales well to high dimensional problems, (b) makes as few function evaluations as possible, and (c) finds a global optimum, given enough time. Unfortunately, none of the existing DFO techniques satisfy all these criteria. Gaussian process optimization [7], perhaps the most popular DFO technique, doesn't scale beyond 100 dimensional problems. In future, we aim to develop DFO techniques that satisfy the above desiderata. We believe our previous work on online nonconvex learning will aid us in designing such algorithms.

2 Boosting

Boosting is a widely used learning technique in machine learning for solving classification problems. Over the years, boosting based methods have shown tremendous success in many real-world applications. Moreover, boosting based methods are easy to train and understand from a theoretical standpoint, thus making it easier to adopt these methods in critical applications such as healthcare. However, this success is mostly limited to classification tasks involving structured or tabular data with hand-engineered features. On classification problems involving low-level features and complex decision boundaries, boosting tends to perform poorly. One example where this is evident is the image classification task, where the decision boundaries are often complex and the features are low-level pixel intensities. This drawback stems from the fact that boosting builds an additive model of weak classifiers, each of which has very little predictive power. Since such additive models with any reasonable number of weak classifiers are usually not powerful enough to approximate complex decision boundaries, the models' output by boosting tend to have poor performance. This then brings us to the following question: *can we generalize boosting to allow for more complex forms of aggregation than linear combinations of weak classifiers?* Such a generalized boosting algorithm can have several benefits. For example, if we can develop boosting algorithms which combine weak classifiers through function compositions, it entails a simple and easy to understand algorithm for learning neural networks. Moreover, such an algorithm can make **neural network training transparent** and easy to adopt in critical applications.

The above question can be studied from two different perspectives: one based on the statistical view of boosting, where boosting is viewed as greedy stagewise optimization [8], and the other based on the game-theoretic view, where boosting algorithms are viewed as playing a game against a weak learner [9].

Generalized boosting from statistical perspective (previous work). In [4], we studied the above question from the statistical viewpoint. In particular, we developed greedy stagewise optimization algorithms which allow for more complex forms of aggregation than additive combinations that are considered by traditional stagewise optimization techniques. Our algorithms improve upon traditional boosting and bridge the gap in performance between traditional boosting and neural networks (see table on the right).

Algorithm	Test Accuracy		
	SVHN	FashionMNIST	CIFAR10
XGBoost (Trees)	77.72	90.34	58.34
AdaBoost (1 layer NNs)	82.88	88	72.78
Generalized Boosting (function compositions)	91.03	93.17	82.31
ResNet	94.82	93.49	86.88

Generalized boosting from game-theoretic perspective (future work). The algorithms we developed in [4] don't yet match the performance of end-to-end trained neural networks. To truly bridge the gap in performance between boosting and neural networks, we hypothesize that one has to look at the game-theoretic viewpoint of boosting. Historically, the game theoretic perspective has been much more successful in developing boosting algorithms with good generalization guarantees, than the statistical perspective. For example, consider the problem of multiclass boosting. Numerous boosting algorithms have been developed for this problem from the statistical perspective. However, many of these algorithms often perform poorly in practice. When viewed from a game-theoretic perspective, many of these algorithms actually turn out to be sub-optimal [10]. Furthermore, the game-theoretic viewpoint has played a crucial role in designing optimal algorithms for multiclass boosting. Consequently, in future, we aim to develop generalized boosting algorithms from a game-theoretic perspective.

3 Research Goals

My research contributes to algorithmic and methodological advances in game theory, online learning, and optimization while addressing classical and emerging problems in statistics and machine learning. I would like to describe my research agenda by stating a few short-term goals and a long-term vision that can set the stage for years of work.

Short term Plans. In the short and medium-term, I am interested in scaling up my work on minimax statistical estimation to high dimensional problems, with the ultimate goal of building an AI/ML system that can aid statisticians in designing optimal estimators for various statistical problems. To this end, I would like to work on developing fast derivative-free optimization methods and scalable Bayesian estimation techniques. I'm also interested in applying the algorithmic tools I develop to emerging game-theoretic problems in machine learning such as robust learning, adversarial density estimation, and to other online learning problems such as contextual bandits. Finally, I plan to continue my work on developing a generalized boosting framework that can match the performance of deep neural networks.

Long term Plans. In the long term, I would like to take a broader perspective and study other aspects of the interplay between game theory and machine learning. My Ph.D. work has mostly focused on studying two-player games in machine learning and statistics. In the future, I plan to move beyond two-player games. Several emerging problems in machine learning naturally lead to *multi-player* games and *sequential* games. Due to various privacy concerns, data used in many modern applications of machine learning in health-care and advertising, is often collected in a decentralized manner by multiple local actors, each with their own self-interests. Statistical inference in such scenarios leads to an interplay with multi-player game theory.

Similarly, sequential games naturally arise in estimation problems where data arrives sequentially, and in problems such as *reinforcement learning* which involve sequential prediction. I plan to work on various aspects of these problems and develop algorithmic and analytic tools in game theory to study them.

References

[1] Kartik Gupta, Arun Sai Suggala, Adarsh Prasad, Praneeth Netrapalli, and Pradeep Ravikumar. Learning minimax estimators via online learning. *arXiv preprint arXiv:2006.11430*, 2020.

[2] Arun Sai Suggala and Praneeth Netrapalli. Online non-convex learning: Following the perturbed leader is optimal. volume 117 of *Proceedings of Machine Learning Research*, pages 845–861, San Diego, California, USA, 08 Feb–11 Feb 2020. PMLR.

[3] Arun Sai Suggala and Praneeth Netrapalli. Follow the perturbed leader: Optimism and fast parallel algorithms for smooth minimax games. In *Advances in Neural Information Processing Systems 33*. 2020.

[4] Arun Sai Suggala, Bingbin Liu, and Pradeep Ravikumar. Generalized boosting. In *Advances in Neural Information Processing Systems 33*. 2020.

[5] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008.

[6] Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, 2015.

[7] Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

[8] Jerome Friedman, Trevor Hastie, Robert Tibshirani, et al. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The annals of statistics*, 28(2):337–407, 2000.

[9] Yoav Freund and Robert E Schapire. A desicion-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer, 1995.

[10] Indraneel Mukherjee and Robert E Schapire. A theory of multiclass boosting. *Journal of Machine Learning Research*, 14(Feb):437–497, 2013.