Research Statement

Algorithmic Mechanism design and game theory are my primary research interests. Mechanism design is an important game theoretic tool in microeconomics. It has found widespread applications in modeling and solving decentralized design problems in many branches of engineering, notably computer science, electronic commerce, and network economics. Mechanism design is concerned with settings where a social planner faces the problem of aggregating the announced preferences of multiple agents into a collective decision when the agents exhibit strategic behaviour [1]. The system designer's goal is to achieve certain desired optimality based on the reported preferences. The strategic agents can manipulate underlying optimization algorithm by misreporting their preferences. Mechanism design is reverse engineering of game theory. It involves inducing a game among agents such that, at equilibrium, the strategic agents report their preferences truthfully.

My research aims to build solutions to mechanism design problems that are having strong theoretical foundation, are of practical use and are computationally efficient. Mechanism design has numerous applications in computer science, esp building autonomous agent systems where the agents are intelligent and strategic. Algorithm design plays important role as the underlying problem that a mechanism solves may be computationally hard. The traditional approximation algorithms might fail to ensure game theoretic properties. Thus, the mechanism needs to design algorithm ingeniously that safeguards incentive constraints.

Allocation of heterogeneous objects or resources to competing agents is a ubiquitous problem in the real world. For example, a federal government may wish to allocate different types of spectrum licenses to telecom service providers; a search engine has to assign different sponsored slots to the ads of advertisers; etc. The agents involved in such situations have private preferences over the allocations. The agents, being strategic, may manipulate the allocation procedure to get a favorable allocation. If the objects to be allocated are heterogeneous (rather than homogeneous), the information that agents posses, needs to be represented in multi-dimensional spaces. This makes the allocation problem complex as the agents get more degrees of freedom for manipulating the algorithm. It becomes even more formidable in the presence of a dynamic supply and/or demand. My doctoral thesis work is motivated by such problems involving strategic agents, heterogeneous objects, and dynamic supply and/or demand. My thesis proposed novel solutions to the problems in resource allocations using techniques from mechanism design theory, game theory, online algorithms and machine learning.

Currently, I am excited about exploring other frontiers of the mechanism design theory namely, algorithmic mechanism design, computational social choice, interface between cryptography and game theory, using game theory and mechanism design for social networks, crowdsourcing. All these problems are mechanism design problems The following are the research directions I would like to pursue in coming years.

Game Theory and Cryptography

Many times participating agents are very sensitive about their preferences and it is desired to take decision without revealing their preferences to the designer [2, 3]. For example, it is very important to design good auctions that preserve the privacy of the bids as well as the fact of bidding itself from the remaining agents. The theory of secure multi-party computation using homomorphic encryption is handy in designing such auction mechanisms [4]. However, the state of the art does not address preserving privacy of the fact of bidding itself in combinatorial auctions, that is who is bidding for what. We have made progress in designing such privacy preserving combinatorial auctions. There are many interesting challenges to be addressed here. Yet most of the proposed solutions having good privacy properties are not scalable and when adopted for combinatorial auctions where underlying auctions themselves are hard, the privacy concerns makes solutions further away from usable. It is desired to build scalable solutions for privacy preserving combinatorial auctions. Privacy preserving exchanges are not yet designed. A distributed constraint optimization with privacy guarantees [5] can be used to design such exchange protocols. I am interested in pursuing these research problems.

Differential Privacy and Mechanism Design for Secure Auctions

The researchers de-anonymized the famous Netflix challenge database, by linking it with IMDb database. Thus, leading to a compromise on individual privacy. The notion of differential privacy, introduced by [6], is very important to address this issue. It protects the data from identifying any specific record while performing statistical queries onto it. The connection between differential privacy and mechanism design was shown in [7, 8]. However, it has not yet been explored to build stronger secure auction protocols. I am interested in building secure auctions using differential privacy.

Learning Mechanisms

Whenever autonomous agents as well as the system designer are learning about the environment to optimize certain parameters whose values are unknown, techniques from machine learning/reinforcement learning are deployed. In such systems, the selfish agents may manipulate the learning algorithms [9, 10, 11] and mechanism design is a natural tool. This calls for combining mechanism design and machine learning techniques, which is referred to as Incentive Compatible Machine Learning. The following examples depict the need of incentive compatible machine learning.

- In sponsored search auctions over search engines like google, an advertiser pays to the search engine only when his advertisement receives a click. This probability of click for the advertiser, popularly known as click through rate (CTR), is an important parameter in the sponsored search auctions. These CTRs are not known beforehand. CTRs can be learned over repeated auctions [9, 12, 10]. This represents a strong case to combine machine learning with mechanism design theory. We have characterized truthful mechanisms in such contexts.

- In the later half of last decade, crowdsourcing has emerged as a new paradigm in outsourcing. Crowdsourcing, according to Jeff Howe, is the act of a company or institution taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call. A challenge here is that the workers are undefined and we do not know the quality of the workers as well as costs of performing the tasks are private to the workers. These qualities can be learned over a period. At the same time, the requesters have to balance between retaining good workers by appropriate incentives and not overpaying spammers. Mechanism design/game theory and machine learning/reinforcement learning need to be combined to design better systems in such contexts. [13, 11] make progress in this direction.

There are many interesting research problems which I am currently focusing on, a complete characterization of truthful machine learning algorithms for general settings is not available. The optimization problems which the system designer is solving may be NP complete and needs to solve it approximately. We need to then design ingenious algorithms which satisfy certain monotonicity for ensuring game theoretic properties. In certain settings, the parameters that mechanism is trying to learn might change over the period. Designing an incentive compatible mechanism with such dynamic nature of the unknown parameters makes the problem more challenging.

Algorithmic Mechanism Design in Multi-dimensional Settings

In his seminal work, Myerson [14] characterized an optimal auction for selling a single unit of a single item. Some limited generalizations for multi-unit single item auctions have been proposed [15, 16]. In multi-item settings, the representation of agents' preferences is multi-dimensional and this makes the mechanism design highly non-trivial. Ensuring truthful behaviour becomes extremely challenging mathematical problem and very little is known for general settings. My doctoral work contributed towards designing mechanisms in special multi-dimensional settings. Here, structure of the problems was exploited to design novel mechanisms. We made progress in designing optimal multi-unit combinatorial auctions under certain assumptions [17]. Until the recent break-throughs [18, 19], exact optimal auctions for multi-items were not available. In [18], designing an optimal multi-item auction is reduced to a problem of social welfare maximization via mapping the preferences of the agents to virtual preferences and then solved using geometric algorithms. However, structure of these transformations for general settings are yet to be analyzed. At the same time, the underlying algorithms for social welfare maximization are computationally too intensive for many interesting situations. Therefore we need to solve such problems approximately but at the same time, we need to ensure that the game theoretic properties of mechanisms are within tolerable limits. This calls for algorithmic mechanism design [20]. I am very much interested in addressing design of computationally efficient, multi-dimensional optimal auction that has good approximation bounds.

Dynamic Mechanism Design without Money

In certain settings, monetary transfers are not allowed. Kidney exchanges, campus recruitment and university dorm allocations are some such scenarios. In these problems the agents have preferences over their matches. Such problems are broadly classified into two categories, one sided matching and two sided matching. The theory developed to address incentive constraints in such settings assumes that all the agents are simultaneously available and they are static. However, in reality the agents are dynamic and are available only for a limited period. The mechanism needs to take a decision pertaining an agent before he or she leaves the system. [21, 22, 23] show that we can adopt static solutions to build dynamic solutions preserving an important game theoretic property, incentive compatibility. However, these solutions lead to poor performance on other desirable properties such as rank efficiency and stability (where no pair of agents can prefer to be matched with each other than their match). There are many interesting unsolved questions here: characterize all mechanisms that can be implemented truthfully, study the right trade-off between incentive constraints and the other desirable properties needs to be studied. I am keen on continuing the work in this direction. In particular, to introduce weaker notion of desirable properties like stability, to look for weaker solution concepts such as Bayesian incentive compatibility and improve on the performance for rank efficiency, stability.

Applications

The research topics discussed above are different sub-fields under the broad umbrella of mechanism design theory. These advances have widespread applications across multiple domains: social networks, crowdsourcing, online advertising, market prediction, intelligent transportation and even online education, to name a few. The participants in such marketplaces are intelligent and rational agents. Therefore, to achieve a large system-wide goal, it is important to provide them with good motivation in the form of proper incentives. Using mechanism design theory, I would like to contribute in building applications for such marketplaces ensuring that the marketplace is robust to agent's manipulations.

References

[1] D. Garg, Y. Narahari, and S. Gujar, "Foundations of mechanism design: A tutorial - part 1: Key concepts and classical results," *Sadhana - Indian Academy Proceedings in Engineering Sciences*, vol. 33, pp. 83–130, April 2008.

[2] A. C. Yao, "Protocols for secure computations," in *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, pp. 160–164, 1982.

[3] S. Izmalkov, S. Micali, and M. Lepinski, "Rational secure computation and ideal mechanism design," in *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pp. 585–594, IEEE, 2005.

[4] F. Brandt and T. Sandholm, "On the existence of unconditionally privacy-preserving auction protocols," *ACM Transactions on Information and System Security (TISSEC)*, vol. 11, no. 2, p. 6, 2008.

[5] T. Léauté and B. Faltings, "Protecting privacy through distributed computation in multi-agent decision making," *Journal of Artificial Intelligence Research*, vol. 47, no. 1, pp. 649–695, 2013.

[6] C. Dwork, "Differential privacy," in *Automata, languages and programming*, pp. 1–12, Springer, 2006.

[7] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pp. 94–103, IEEE, 2007.

[8] K. Nissim, R. Smorodinsky, and M. Tennenholtz, "Approximately optimal mechanism design via differential privacy," in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 203–213, ACM, 2012.

[9] M. Babaioff, Y. Sharma, and A. Slivkins, "Characterizing truthful multi-armed bandit mechanisms," in *EC09: Proceedings of the 10th ACM Conference on Electronic Commerce*, (Stanford, California), pp. 79–88, 2009.

[10] A. Das Sharma, S. Gujar, and Y. Narahari, "Truthful multi-armed bandit mechanisms for multi-slot sponsored search auctions," *CURRENT SCIENCE*, vol. 103, no. 9, pp. 1064–1077, 2012.

[11] S. Jain, S. Gujar, O. Zoeter, and Y. Narahari, "A quality assuring multi-armed bandit crowdsourcing mechanism with incentive compatible learning," in *Thirtheenth International Conference on Autonomous Agents and Multiagent Systems*, pp. 1609–1610, 2014.

[12] N. R. Devanur and S. M. Kakade, "The price of truthfulness for pay-per-click auctions," in *EC09: Proceedings of the 10th ACM Conference on Electronic Commerce*, pp. 99–106, 2009.

[13] S. Bhat, S. Nath, O. Zoeter, S. Gujar, Y. Narahari, and C. Dance, "A mechanism to optimally balance cost and quality of labeling tasks outsourced to strategic agents," in *Thirtheenth International Conference on Autonomous Agents and Multiagent Systems*, pp. 917–924, 2014.

[14] R. B. Myerson, "Optimal auction design," *Mathematics of Operations Research*, vol. 6, pp. 58–73, February 1981.

[15] G. Iyengar and A. Kumar, "Optimal procurement mechanisms for divisible goods with capacitated suppliers," *Review of Economic Design*, vol. 12, pp. 129–154, June 2008.

[16] A. Malakhov and R. V. Vohra, "Single and multi-dimensional optimal auctions - a network perspective," discussion papers, Kellogg School of Management, Northwestern University, December 2005.

[17] S. Gujar and Y. Narahari, "Optimal multi-unit combinatorial auctions," *Operational Research*, vol. 13, pp. 27–46, 2013.

[18] Y. Cai, C. Daskalakis, and S. M. Weinberg, "Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization," in *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pp. 130–139, IEEE, 2012.

[19] Y. Cai, C. Daskalakis, and S. M. Weinberg, "An algorithmic characterization of multi-dimensional mechanisms," in *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 459–478, ACM, 2012.

[20] N. Nisan and A. Ronen, "Algorithmic mechanism design," in *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pp. 129–140, ACM, 1999.

[21] S. Gujar and D. C. Parkes, "Dynamic Matching with a Fall-back Option," in *In Proceedings of 19th European Conference on Artifitial Intelligence*, pp. 263–268, 2010.

[22] J. Zou, S. Gujar, and D. C. Parkes, "Tolerable Manipulability in Dynamic Assignment without Money," in *Proceedings 24th AAAI Conference on Artificial Intelligence (AAAI '10)*, 2010.

[23] S. Guajr and B. Faltings, "Dynamic task assignments: An online two sided matching approach," in *Proceedings of the 3rd International workshop on Matching Under Preferences, MATCHUP, to appear*, 2015.