Research Statement

Michael Munie

In recent years there has been growing interest in the interaction between computer science and the social sciences (in particular, social-choice theory and game theory), spurred by electronic commerce and other aspect of the internet. My work falls into this emerging area, though with a slightly different emphasis than much of the work in the area. Specifically, my work focuses on the dynamics and convergence of online systems where selfish agents interact with an aggregating center.

Previous Work

My past research has been on two areas, aggregation systems and voting systems.

Wikipedia, Yelp, and Trip advisor are all examples of aggregation systems. These sites are all online processes that take input from history-aware users and display some aggregation of these inputs. On Yelp, users choose to give a restaurant one to five stars and a summary rating is displayed. How will users choose their rating? Will the ratings of a location converge over time, or will users cause the summary rating to continually fluctuate? Furthermore, what does the aggregate output mean? Is it some statistic of user's possibly hidden information given that they may not take actions that accurately express their internal beliefs?

I proposed a new class of games [1, 2], historical games, to study this problem. They are infinite games, but where the current round game is determined by the history of all past actions. Traditional game theory tools like Nash Equilibrium and Folk Theorems still apply, but by also looking at myopic action selection rules we prove that the play and payoff in these games always converge. The convergence point is a median of agent's goals and social focal points defined by the relative influence of agents. This convergence point is close to the mean of the agent's private information, but not the same.

Furthermore, historical games are closely connected to social choice theory. A special class of historical games in single-peaked domains has a direct correspondence between every strategy-proof and efficient social choice function. In this class of voting rules, agents reveal their preferences to the center, but while agents acting in historical games converge to the same point they reveal much less information.

The behavior of one of these systems can be seen in Figure 1. By first approaching this problem through large scale simulation, I was able to gain insight into the dynamics, and finally formally prove the convergence results and correspondence to strategy-proof voting.

The second area I've focused on is voting systems [3]. My research turns the problem of social choice on its head as an optimization problem. I take the traditional axioms and view axiomatic violations as errors. Based on these violations, I propose new metrics for the quality of voting rules. In this research I compare voting rules by computationally estimating the probability that an election will produce an error, and prove analytically that these probabilities converge to zero for all scoring rules. In addition, I also defined a new data structure, the violation graph, for computing these metrics and as visualization for comparing the nature of axiomatic violations. In Figure 2 four common voting rules are represented as violation graphs (only connected nodes are shown) where nodes represent different preference profiles for three voters over three candidates. In the small example we can see that not only do the number of violations as measured by nodes or edges vary greatly, but also some voting rules (copeland) need only a small number of nodes removed before there are no violations at all.

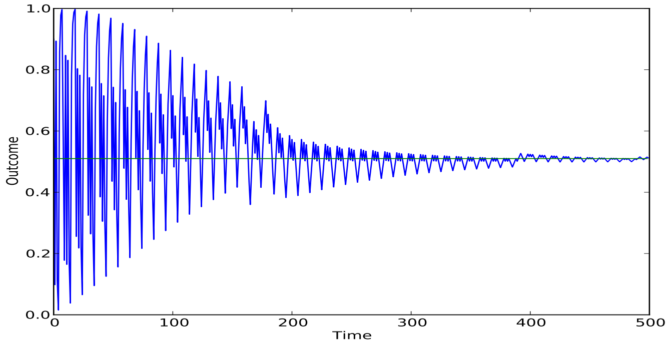

Figure 1: Convergence under myopic play

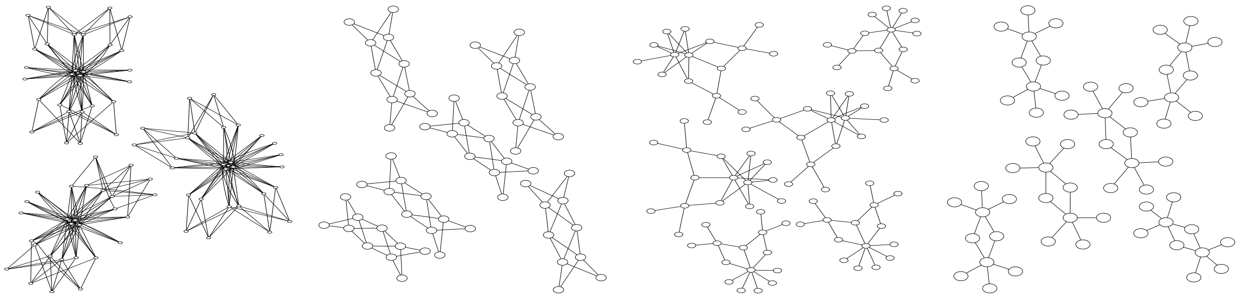

(a) Plurality (b) Plurality with Runoff (c) Borda (d) Copeland

Figure 2: Violation Graphs for common voting rules

Future Work

My vision for my future work is to understand the dynamics and design of systems where selfish agents continue to interact over time and influence the nature of their future interactions. Specifically, this means first proving a wider class of systems convergence, then focusing on the design of systems to achieve desired converge points, and finally studying the design of organizational structure.

In the short term, I plan to focus on applying my work on aggregation systems to higher dimensional preference spaces. Research from analytical political science has shown that with both separable and non-separable preferences, voting in higher dimensional preference spaces does not converge. By running large scale simulations I have evidence that in historical games, even with non-separable preferences, actions and payoffs converge. My goal is to extend my work on convergence in one dimension to convergence in higher dimensions which more accurately model many real world problems.

In the longer term, I want to work on the larger problem of computational society design – how do we structure societies for good decisions? This can be broken down into two lines of work.

The first is aggregation design. A superset of my previous work, the larger question is what should an aggregation environment where multiple self-interested agents are interacting with a system achieve? What should the aggregate tell us: the mean, variance, or something else? Is convergence necessary or desirable? The agents can interact with the system in different ways and as a designer it is necessary to consider the reasoning methods of the agents along with their goals in order to achieve the desired outcome.

In other environments agents may directly communicate with each other and we may need to delegate powers and responsibilities. Here we ask: how should agents be organized to jointly achieve a goal in a changing environment? This task is called organization design. For example, how should we structure a corporation to maximize average output, but minimize worst case output? Communication between agents, computational powers of individual agents, and control between agents are all factors must be considered. Ideally, we would be able to approach this problem from an algorithmic viewpoint and given an environment, output the optimal structure.

My research is theoretically grounded but motivated by empirical results from large scale simulations. It ties together political science, economics, and computer science. With the talented researchers and deep knowledge of weaving together economics and computer science, Microsoft research is the best place for me to do this work.

References

[1] M. Munie, Y. Shoham, "Historical Games", *Under Submission*.

[2] M. Munie, Y. Shoham, "Joint Process Games: From Ratings to Wikis", *In Proceedings of 9th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS10)*, 2010, Toronto, Canada

[3] M. Munie, P. Tang, and Y. Shoham "A framework for the Quantitative Evaluation of Voting Rules", *In Proceedings of Logic, Game Theory, and Social Choice (LGS6)*, 2010, 2009, Tsukuba, Japan

[4] M. Munie, Y. Shoham, "Optimal Testing of Structured Knowledge", *In Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-08)*, 2008, Chicago, IL