

Unleashing the Power of Multi-Agent Voting Teams

Leandro Soriano Marcolino

Supervisor: Milind Tambe

University of Southern California, Los Angeles, CA, USA

sorianom@usc.edu

Abstract

Teams of voting agents have great potential in finding optimal solutions. However, there are fundamental challenges to effectively use such teams: (i) selecting agents; (ii) aggregating opinions; (iii) assessing performance. I address all these challenges, with theoretical and experimental contributions.

1 Introduction

Teams of voting agents are used in important domains, such as: machine learning, crowdsourcing, forecasting systems, and even board games. Voting is popular since it is highly parallelizable, easy to implement and provide theoretical guarantees [Conitzer and Sandholm, 2005]. However, there are fundamental challenges: (i) Selecting a limited number of agents; (ii) Combining the opinions of the team members; (iii) Assessing the performance. In my thesis, I address all these challenges, with both theoretical and experimental contributions. I explore three different domains: Computer Go, HIV prevention in social networks and architectural design.

Concerning agent selection, I study the importance of diversity. While by previous works in social choice we would expect the best teams to be uniform teams composed of copies of the best agent [Conitzer and Sandholm, 2005], I show in Marcolino *et al.* [2013] the importance of considering diversity when forming teams. My first model only gave necessary conditions, however, so in Marcolino *et al.* [2014] I create a second new model that predicts that diverse teams are better than uniform teams in problems with a large action space.

Concerning aggregating opinions, I study different ranked voting rules and ranking extraction techniques. Recently, ranked voting has received considerable attention [Caragiannis *et al.*, 2013]. However, I show that plurality still outperforms ranked voting rules in the Computer Go domain [Jiang *et al.*, 2014], and in the HIV prevention domain [Yadav *et al.*, 2015]. This is caused by the noise in the rankings of agents that were originally designed to output a single choice. Therefore, in Jiang *et al.* [2014], I introduce a new ranking extraction technique, based on the frequency that actions are played when sampling an agent multiple times.

Concerning assessing performance, I introduce a novel domain-independent technique that allows one to predict whether a team of voting agents is going to be successful or

fail in problem solving. Such prediction is important to take remedy procedures to increase a team's performance. Existing methods are tailored for specific domains [Ramos and Ayanegui, 2008]. In Nagarajan *et al.* [2015] I introduce a novel domain independent technique, which learns a prediction function using only the voting patterns of a team.

Finally, I am currently exploring domains where the objective is not simply to find one optimal action, but rather to maximize the number of optimal solutions found. This is useful in design, where a human can take an aesthetical choice among all optimal solutions. In Marcolino *et al.* [2015] I perform a theoretical study of how diverse an uniform teams perform for such problems, and I present experimental results in the architectural building design domain.

2 Theoretical Results

I briefly summarize some of my theoretical models. The first, presented in Marcolino *et al.* [2013], shows that a diverse team can outperform a uniform team if at least one agent has a higher probability of playing the best action than the best agent in at least one world state. Since only necessary conditions were provided, I study in Marcolino *et al.* [2014] the effect of increasing the number of actions available to choose from. I define *spreading tail (ST)* agents, that have an increasingly larger number of actions assigned with a non-zero probability as the number of actions in the domain increases. A diverse team is modeled as a team of *ST* agents. I show that the probability of a diverse team picking the best action increases as the action space increases; and also converges to 1 when the team size grows in large action spaces. The main idea is that diverse agents are less likely to agree on the same mistakes when the action space is large, and therefore only two agents voting for the optimal solution is sufficient.

In Nagarajan *et al.* [2015], I propose a model for team assessment, which allows me to develop a novel domain independent technique to predict online the final reward of a team. The final reward is defined by a random variable; which is influenced by a set of variables \mathbf{H}_j representing the subset of agents that agreed on the chosen action at world state j . Based on that, I show that the final reward can be predicted by linear models, and I derive a prediction function using the frequencies of agreement of each subset of agents.

Finally, I am currently working in a model of agent teams for design problems [Marcolino *et al.*, 2015], where the ob-

jective is to find as many optimal solutions as possible across multiple voting iterations. I show that, unless the best agent has the same probability of voting for all optimal actions, a uniform team will invariably converge to a single optimal solution. On the other hand, a diverse team composed of agents with different preferences maximizes the number of optimal solutions, given some conditions on the team size.

3 Experimental Results

I summarize here some of the experimental results. Figure 1 shows one of the results in Computer Go. We see the winning rates of a *diverse* (composed by the Computer Go playing agents Fuego, Gnugo, Pachi, Mogo) and a *uniform* team. The *diverse* team starts by playing slightly worse, but plays better than the *uniform* team with statistical significance in large boards [Marcolino *et al.*, 2014].

Figure 2 shows one result from my novel ranking methodology, where a ranking is built by sampling each agent 10 times. All tested voting rules outperform plurality, but in the figure we only show Borda (which is better with $p < 0.007$). All voting rules are also statistically significantly better than the non-sampled (single run) version of plurality [Jiang *et al.*, 2014].

Concerning the team assessment problem, Figure 3 shows the accuracy of my predictions for the *diverse*, the *uniform* and an *intermediate* team (composed by different parameterizations of a single agent). The prediction accuracy is close to 60% in the middle game, and it goes to 73% towards the end of the games [Nagarajan *et al.*, 2015].

Figure 4 shows some of the results in the building design domain, where we see the percentage of optimal solutions found by individual agents and by different teams, for three different building design problems. As we can see, the teams clearly outperform the individual agents, and provide a higher percentage of optimal solutions for a designer to choose from [Marcolino *et al.*, 2015].




Figure 1: Winning rates as the Go board size grows.




Figure 2: New ranking method results.




Figure 3: Accuracy for 3 different teams.



Figure 4: Percentage found by each system.

Finally, Figure 5 shows one result in the HIV prevention domain, when choosing which nodes of a social network should be selected across multiple iterations, in order to maximize the influence of HIV prevention techniques. DC and POMCP are baselines, while PSINET-S, PSINET-W and PSINET-C are teams using simple plurality, weighted plurality and Copeland ranked voting, respectively. As we can see, weighted plurality has the best result in this domain [Yadav *et al.*, 2015].




Figure 5: Influence in social networks.

4 Next Steps

There are many open directions for further research. I am currently studying the performance of mixed voting teams of humans and artificial agents. Moreover, concerning the work in HIV prevention [Yadav *et al.*, 2015], it is still open to understand why the Copeland ranked voting rule did not perform well, and how can the performance be improved with novel aggregation and/or ranking extraction methodologies.

Acknowledgments: This research was supported by MURI grant W911NF-11-1-0332.

References

[Caragiannis *et al.*, 2013] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *EC*, 2013.

[Conitzer and Sandholm, 2005] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *UAI*, 2005.

[Jiang *et al.*, 2014] A. X. Jiang, L. S. Marcolino, A. D. Procaccia, T. Sandholm, N. Shah, and M. Tambe. Diverse randomized agents vote to win. In *NIPS*, 2014.

[Marcolino *et al.*, 2013] L. S. Marcolino, A. X. Jiang, and M. Tambe. Multi-agent team formation: Diversity beats strength? In *IJCAI*, 2013.

[Marcolino *et al.*, 2014] L. S. Marcolino, H. Xu, A. X. Jiang, M. Tambe, and E. Bowring. Give a hard problem to a diverse team: Exploring large action spaces. In *AAAI*, 2014.

[Marcolino *et al.*, 2015] L. S. Marcolino, H. Xu, D. Gerber, B. Koley, S. Price, E. Pantazis, and M. Tambe. Agent teams for design problems. In *COIN Workshop*, 2015.

[Nagarajan *et al.*, 2015] V. Nagarajan, L. S. Marcolino, and M. Tambe. Every team deserves a second chance: Identifying when things go wrong. In *AAMAS*, 2015.

[Ramos and Ayanegui, 2008] F. Ramos and H. Ayanegui. Discovering tactical behavior patterns supported by topological structures in soccer-agent domains. In *AAMAS*, 2008.

[Yadav *et al.*, 2015] A. Yadav, L. S. Marcolino, E. Rice, R. Petering, H. Winetrobe, H. Rhoades, M. Tambe, and H. Carmichael. Preventing HIV spread in homeless populations using PSINET. In *AAAI*, 2015.