

Three Fundamental Pillars of Multi-agent Team Formation (Doctoral Consortium)

Leandro Soriano Marcolino
Supervisor: Milind Tambe
University of Southern California
Los Angeles, CA, USA
sorianom@usc.edu

ABSTRACT

Teams of voting agents are a powerful tool for solving complex problems. When forming such teams, there are three fundamental issues that must be addressed: (i) Selecting which agents should form a team; (ii) Aggregating the opinions of the agents; (iii) Assessing the performance of a team. In this thesis we address all these points.

Categories and Subject Descriptors

I.2.1 [Artificial Intelligence]: Applications and Expert Systems

Keywords

Teamwork; Collective intelligence; Social choice theory

1. INTRODUCTION

Teams of voting agents are a powerful tool for solving complex problems in many important domains: crowdsourcing, board games, machine learning, forecasting systems, etc. However, it is challenging to form such teams. I take the position that team formation is not only selecting the team members, but also includes the study of techniques to combine their opinions, and to assess their performance in order to perform readjustments. Therefore, I study three (complementary) facets of team formation: *Agent Selection*, *Aggregation of Opinions* and *Team Assessment*.

Agent Selection is picking a limited number of agents to form a team. When considering previous work in social choice, we would expect the best team to be the one composed of the best possible agents [2]. I show in [5] that this is not true, and it is fundamental to also consider diversity when forming teams. However, [5] only presents necessary conditions for such phenomenon; therefore in [6] I present a more general model of diversity, where I can predict that diverse teams perform better than uniform teams (i.e., copies of the best agent) in problems with a large action space.

Aggregation of Opinions is combining the opinions of each member of the team into a final team decision. Although recently ranked voting has received considerable attention in the social choice literature [1], I show in [3] that plurality

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.
Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

clearly outperforms ranked voting approaches in an actual domain (Computer Go). Such surprising result may seem discouraging at first, but I verify that it is caused by the noise in the rankings of agents that were originally designed to output a single choice. In order to solve this problem, I introduce a novel method to extract a ranking from existing agents, based on the frequency that actions are played when sampling an agent multiple times.

Team Assessment is verifying the performance of a given team. In particular, it is fundamental to predict whether a team is going to be successful or fail in problem solving. Existing methods are tailored for specific domains, such as robot-soccer [8]. Hence, in [7], I propose a novel *domain independent* technique, which learns a prediction function for estimating the performance of a team using only its voting patterns. Based on such prediction, it is possible to take remedy procedures to increase a team's performance.

My thesis presents theoretical and experimental contributions for these different facets of team formation. Moreover, I study three domains: Computer Go [5, 6, 3, 7], Building Design [4] and HIV prevention in social networks [9].

2. THEORETICAL RESULTS

In this section I briefly summarize some of the models developed. The first two focus in the agent selection process. In the first model, each agent is modeled by a probability distribution function $\mathbf{V}_{i,j}(r)$, that gives the probability of agent i voting for an action with rank r (for example, $r = 1$ gives the best action) in world state j . A uniform team is modeled as repeated samples from the same pdf, while a diverse team samples different pdfs. In [5], I show 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.

The second model still represents the agents by pdfs, but focuses in a single world state [6]. I study what happens as the number of actions available to choose from gets larger. 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 study the convergence of the probability of a diverse team picking the best action, and show that it converges to a certain value that is larger than any other value obtained in lower action space sizes. Moreover, I show that if the action space is large enough, the probability of a diverse team choosing the best action converges to 1 as the number of agents increases.

I also develop models for team assessment. In [7], I de-

velop a model where the final reward of a team is defined by a random variable; which is influenced by a set of variables \mathbf{H}_j that represents the subset of agents that agreed on the chosen action at each world state j . I consider that all world states equally influence the outcome, and they are independent. Based on this model, I show that the final outcome of a team can be predicted by linear models, and I also derive a domain independent prediction function using only the frequencies of agreement of each subset of agents.

3. EXPERIMENTAL RESULTS

I summarize here some of the main experimental results obtained so far, in the Computer Go domain and in the building design domain. In Figure 1, we see one of the results in Computer Go. We can see the winning rates of a *diverse* (composed of the Computer Go playing agents Fuego, Gnugo, Pachi, Mogo) and a *uniform* team. As we can see, the *diverse* team starts by playing slightly worse, but plays better than the *uniform* team with statistical significance in large boards [6].

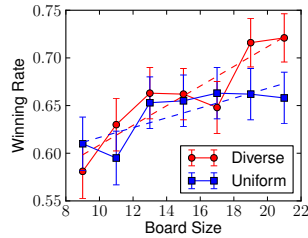


Figure 1: Winning rate in Computer Go, as the board size grows.

In Figure 2, I show one result from my novel ranking methodology. In these results, 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 ranked voting rules are also statistically significantly better than the non-sampled (single run) version of plurality [3].




Figure 2: Borda outperforms plurality, using the new ranking method.

Concerning the team assessment problem, I present in Figure 3 the accuracy of my predictions for the *diverse*, the *uniform* and an *intermediate* team (composed of different parameterizations of a single agent). As can be observed, the prediction accuracy is close to 60% in the middle game, and it goes all the way to 73% towards the end of the games [7].




Figure 3: Accuracy of predictions for 3 different teams of voting agents.

Finally, in Figure 4, I show some of the results in the building design domain. A human designer inputs an initial parametric design of a building, which can be optimized generating a set of optimal design variations. The figure

shows 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. More information at [4].

4. NEXT STEPS

There are still many open directions for further development. I am currently working in a theoretical model of agent teams for design problems, where the objective is to maximize the number of optimal solutions that we can find. Such theory would allow a better understanding of the results in




Figure 4: Percentage of optimal solutions found by each system.

the building design domain. I also expect to develop new techniques for the aggregation of opinions, in order to be able to eventually overcome the performance of the best parallelization algorithm available for Computer Go. It is also possible to use machine learning in order to categorize different kinds of world states (such as positions of the Go board), and based on that information dynamically change the team and/or the voting rule. Hence, many different approaches can still be developed and studied in order to unleash the potential of diverse teams and aggregation of opinions.

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

REFERENCES

- [1] I. Caragiannis, A. D. Procaccia, and N. Shah. When do noisy votes reveal the truth? In *EC*, 2013.
- [2] V. Conitzer and T. Sandholm. Common voting rules as maximum likelihood estimators. In *UAI*, 2005.
- [3] 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.
- [4] L. S. Marcolino, D. Gerber, B. Kolev, S. Price, E. Pantazis, Y. Tian, and M. Tambe. Agents vote for the environment: Designing energy-efficient architecture. In *AAAI Workshop*, 2015.
- [5] L. S. Marcolino, A. X. Jiang, and M. Tambe. Multi-agent team formation: Diversity beats strength? In *IJCAI*, 2013.
- [6] 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.
- [7] V. Nagarajan, L. S. Marcolino, and M. Tambe. Every team deserves a second chance: Identifying when things go wrong. In *AAMAS*, 2015.
- [8] F. Ramos and H. Ayanegui. Discovering tactical behavior patterns supported by topological structures in soccer-agent domains. In *AAMAS*, 2008.
- [9] 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 *IAAI*, 2015.