Multi-agent Team Formation: Solving Complex Problems by Aggregating Opinions

Leandro Soriano Marcolino
University of Southern California, Los Angeles, CA, 90089, USA
sorianom@usc.edu

Introduction

Aggregating the opinions of different agents is a powerful way to find high-quality solutions to complex problems. However, when using agents in this fashion, there are two fundamental open questions. First, given a universe of agents, how to quickly identify which ones should be used to form a team? Second, given a team of agents, what is the best way to aggregate their opinions?

Many researchers value diversity when forming teams. LiCalzi and Surucu (2012) and Hong and Page (2004) propose models where the agents know the utility of the solutions, and the team converges to the best solution found by one of its members. Clearly in complex problems the utility of solutions would not be available, and agents would have to resort to other methods, such as voting, to take a common decision. Lamberson and Page (2012) study diversity in the context of forecasts, where the solutions are represented by real numbers and the team takes the average of the opinion of its members. Domains where the possible solutions are discrete, however, are not captured by such a model.

I proposed a new model to study teams of agents that vote in discrete solution spaces (Marcolino, Jiang, and Tambe 2013), where I show that a diverse team of weaker agents can overcome a uniform team made of copies of the best agent. However, this phenomenon does not always occur, and it is still necessary to identify when we should use diverse teams and when uniform teams would be more appropriate.

Hence, in Marcolino et al. (2014b), I shed a new light into this problem, by presenting a new, more general model of diversity for teams of voting agents. Using that model I can predict that diverse teams perform better than uniform teams in problems with a large action space.

All my predictions are verified in a real system of voting agents, in the Computer Go domain. I show that: (i) a team of diverse players gets a higher winning rate than a uniform team made of copies of the best agent; (ii) the diverse team plays increasingly better as the board size increases.

Moreover, I also performed an experimental study in the building design domain. This is a fundamental domain in the current scenario, since it is known that the design of a building has a major impact in the consumption of energy throughout its whole lifespan (Lin and Gerber 2014). It is fundamental to design energy efficient buildings. Meanwhile, it is important to balance other factors, such as construction cost, creating a multi-objective optimization problem. I show that by aggregating the opinions of a team of agents, a higher number of 1^{st} ranked solutions in the Pareto frontier is found than when using a single agent. Moreover, my approach eliminates falsely reported 1^{st} ranked solutions (Marcolino et al. 2014a; 2015).

As mentioned, studying different aggregation rules is also fundamental. In Jiang et al. (2014), I introduce a novel method to extract a ranking from agents, based on the frequency that actions are played when sampling them multiple times. My method leads to significant improvements in the winning rate in Go games when using the Borda voting rule to aggregate the generated rankings.

Diversity Models

I briefly summarize here the diversity models developed. In the first, 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 Marcolino, Jiang, and Tambe (2013), I show that a diverse team can outperform a uniform team only if at least one agent has a higher probability of playing the best action than the best agent in at least one world state. Moreover, I also show that breaking ties in favor of the best agent is the optimal voting rule for a diverse team. My experimental analysis in the Computer Go domain shows not only that a diverse team actually outperforms a uniform team, but also that weak agents are really able to play better than the best agent for some world state configurations.

The second model still represents the agents by pdfs, but focuses in a single world state (Marcolino et al. 2014b). 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.

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 by 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.

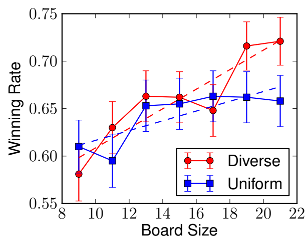

Figure 1: Winning rate in the real Computer Go system, as the board size grows.

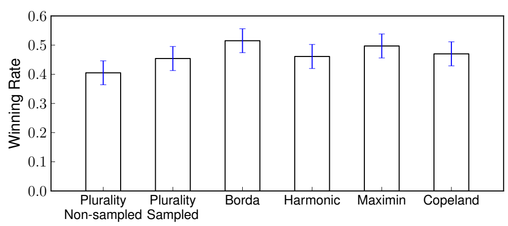

Figure 2: All voting rules, for *diverse* with 5 agents, using the new ranking methodology.

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. We can see that all voting rules outperform plurality; Borda and maximin are statistically significantly better ($p < 0.007$ and $p = 0.06$, respectively). All ranked voting rules are also statistically significantly better than the non-sampled (single run) version of plurality.

Finally, in Figure 3 I show one of the results in the building design domain. The figure shows the number of unique 1^{st} ranked solutions found by individual agents and by different teams, for one building design problem. As we can see, the teams clearly outperform the individual agents, and provide a higher number of 1^{st} ranked solutions.

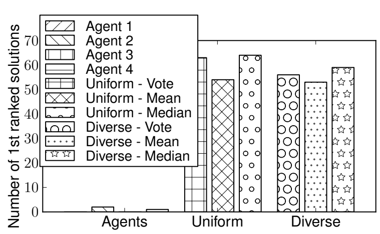

Figure 3: Number of unique 1^{st} ranked solutions, for each agent and the teams.

Next Steps

Many research directions are still open, and I will keep developing further. I plan to defend my thesis in May 2016, so there is still a significant amount of research to be developed.

Moving forward, I expect to perform a theoretical study of diverse teams in multi-objective optimization problems, in order to better understand the results in 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.

Finally, I also recently observed that it is possible to use the voting patterns to predict the final performance of teams of voting agents. Such "side-effect" of voting has not been observed before, and it allows an operator to take remedy procedures if the performance of a team presents issues. I am currently performing a deeper study of this phenomenon.

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

References

Hong, L., and Page, S. E. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proc. of the National Academy of Sciences of the USA* 101(46):16385–16389.

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

Lamberson, P. J., and Page, S. E. 2012. Optimal forecasting groups. *Management Science* 58(4):805–810.

LiCalzi, M., and Surucu, O. 2012. The power of diversity over large solution spaces. *Management Science* 58(7):1408–1421.

Lin, S.-H. E., and Gerber, D. J. 2014. Evolutionary energy performance feedback for design. *Energy and Buildings* 84:426–441.

Marcolino, L. S.; Kolev, B.; Price, S.; Veetil, S. P.; Gerber, D.; Musil, J.; and Tambe, M. 2014a. Aggregating opinions to design energy-efficient buildings. In *M-PREF 2014*.

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

Marcolino, L. S.; Gerber, D.; Kolev, B.; Price, S.; Pantazis, E.; Tian, Y.; and Tambe, M. 2015. So many options, but we need them all: Agent teams for design problems. In *AAMAS*. (submitted).

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