Diversity Beats Strength? A Hands-on Experience with 9x9 Go (Demonstration)

Leandro Soriano Marcolino, Douglass Chen, Albert Xin Jiang, Milind Tambe
University of Southern California, Los Angeles, CA, 90089
{sorianom, douglawc, jiangx, tambe}@usc.edu

ABSTRACT

Team formation is a critical step in deploying a multi-agent team. In some scenarios, agents coordinate by voting continuously. When forming such teams, should we focus on the diversity of the team or on the strength of each member? Can a team of diverse (and weak) agents outperform a uniform team of strong agents? In this demo, the user will be able to explore these questions by playing one of the most challenging board games: Go.

Categories and Subject Descriptors

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

Keywords

Collective Intelligence; Emergent Behavior; Voting

1. INTRODUCTION

Team formation is essential when dealing with a multi-agent system. Given limited resources, we must select a strong team to deal with a complex problem. After forming a team, their members must work together. There are many different ways for a team to coordinate. One common and simple way is to use voting. By voting, a team of agents can get closer to finding the best possible decision in a given situation [5]. Only one voting iteration might not be enough, sometimes the agents must vote continuously in many different scenarios. Consider, for example, agents that are cooperating in a board game [7], deciding together stock purchases across different economic scenarios, or even picking items to recommend to a large number of users [1]. This situation imposes a conflict for team formation: should we focus on the diversity of the team or on the strength of each individual member? Previous works do not address this issue.

Diversity is proposed as an important concept for team formation in the field of Economics and Social Science [3, 4]. However, they assume a model where each agent brings more information, and the system converges to one of the best options known by the group. When a team votes to decide its final opinion, their model and theorems do not hold anymore. It is necessary to develop, therefore, a new model to analyze a team of voting agents.

In [6] we presented a new theory for diversity and strength in teamwork. The fundamental novelty of our model is to consider a setting with multiple world states, and each agent having different performance levels across world states. We presented an extensive experimental analysis on the Computer Go domain, one of the main challenges for Artificial Intelligence. We used 6 heterogeneous Go playing agents, to show that a team of agents performing voting algorithms can play better than each one of its members alone. We also showed that in some situations a team of weak but diverse agents can play better than a team of strong homogeneous agents, and better than a parallelized agent. The fundamental research question we are addressing is understanding the impact of diversity and strength on a team of agents.

In this demo, the user will have the opportunity to explore this question while playing a 9x9 Go game against a team of heterogeneous agents. The agents will decide their action by performing weighted voting, without human intervention. The user will be able to see which teams are stronger and which teams are weaker by her own experience on the board. We will also allow multiple users, and the possibility to play with a team of agents, leading to mixed teams of humans and agents playing against each other. A video of our system can be seen in `http://youtu.be/PswMwnkpsWA`.

2. RESEARCH APPROACH

In generic terms, our team of agents must take iterative decisions in an action space across different world states, and we want to maximize a certain reward in the final iteration. We define diversity as agents having different probability distributions for selecting their actions in different world states. Therefore, in a fixed world state, each agent will have a different expected utility. The strongest agent does not necessarily have the highest expected utility in all world states, even though it plays the best overall. We define strength as the average of the expected utility over all world states.

At each iteration, each agent examines the current world state and submits its (single) opinion about which should be the next action. The opinions are then combined using a weighted majority voting approach. Given a certain weight w for each one of the agents, we sum up the weights for each action a: $a_o = \sum_{i \in N_o} w_i$, where a_o is a certain possible action, N_o is the set of all agents that voted for that action a_o, and w_i is the weight of each agent i.

We then select the action with the highest weight as the next action for that particular iteration. If there is a tie between two or more actions, we pick one of them randomly, according to a uniform distribution. This voting procedure repeats at every iteration, until the end, when the system can obtain a reward.

While demonstrated in the Computer Go domain, the research underlying this demonstration attempts to answer the question of how to select team members: whether to base the team on strength or on diversity. Our model, unlike that in [3, 4], allows us to pro-

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

vide some theoretical guarantees in the context of weighted voting. In our theorems presented in [6], we are able to show that for a team of diverse agents to perform better than a team of non-diverse but strong agents either one of two conditions are necessary: (i) In some world states, some agents in the diverse team must have a higher probability of selecting the best action than the agents in the non-diverse team, even though the agents in the non-diverse team are stronger overall; and/or (ii) In some world states, some agents in the diverse team must have a lower probability of selecting a sub-optimal action than the agents in the non-diverse team. We are also able to show that given some conditions detailed in [6] (intuitively, that all agents contribute to the team) breaking ties in favor of the strongest agent is the optimal voting rule for a diverse team.

For readers familiar with Go, a key question is if the agents will actually agree in some movements or disagree most of the time. This question arises because the 9x9 Go board has many possible movements, reaching 81 options to play in the beginning of the game. If the agents always disagree, the system would be dominated by the agent with the highest weight, or simply randomly choose between agents with the same weight. However, in our experiments we saw that the agents actually only disagree completely in a small number of occasions. The user will be able to observe this phenomenon while playing our demo.

3. DEMONSTRATION

Go is a turn-based game between two players: black and white. At each turn, the players must place a stone in an empty intersection of the board. If a group of stones is surrounded by the opponent's stones they are removed from the board. The stones that surround an area form a territory, whose value is counted by the number of empty intersections inside. In the end of the game, the score is defined by the amount of territory minus the number of captured stones, and the player with the highest score wins. A detailed description of the rules can be found in [8].

Our system has four different Go software: Fuego 1.1, GnuGo 3.8, Pachi 9.01, MoGo 3, and two (weaker) variants of Fuego, in a total of 6 different agents. These are all publicly available Go software. Fuego is known to be the strongest Go software among all of them. Fuego, Pachi and MoGo all follow a UCT Monte Carlo Go algorithm [2].

The user first determines the team that she wants to play against, by dragging agents from an agent pool (Figure 1(a)), and selecting their relative weights (Figure 1(b)). The agents are represented by characters, giving a visual representation of diversity and strength. The user will then play the game in a graphical interface, by clicking in the position where she wants to play. When it is our system's turn, the program displays the votes of all agents in the Go board, followed by the final decision according to the weighted voting rule (Figure 1(c)). The process repeats until the end of the game. In another screen, the system displays statistics about the game, including the frequency that the votes of each agent was accepted and the expected size of the set of agents that voted for the accepted move (Figure 1(d)).

Our demo also have different ways to interact with the system. Besides playing a game against a team of agents, following the previous description, the system also allows:

- The user to be a part of a team of agents, and participate in the weighted voting procedure;
- Multiple users playing simultaneously, leading to mixed teams of humans and agents playing against each other;

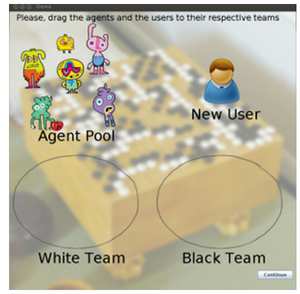

(a) User selects a team of agents.

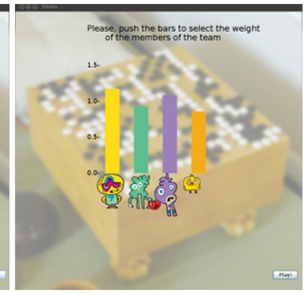

(b) User selects the relative weight of the agents.

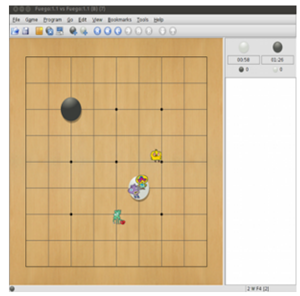
(c) User plays her move. The system replies, while displaying agents' votes.

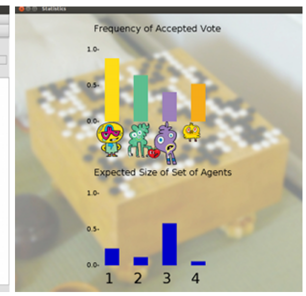

(d) Another screen display statistics about the game.

Figure 1: Example of an execution.

- The user to observe an agent or a team of agents playing against another team.

Therefore, while playing a game, the user will have a hands-on experience in our demo with the decision-making process of a voting-based system for combining a team of diverse cooperative agents.

4. REFERENCES

[1] R. Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, Nov. 2002.

[2] S. Gelly, Y. Wang, R. Munos, and O. Teytaud. Modification of UCT with patterns in Monte-Carlo Go. Technical report, 2006.

[3] L. Hong, S. E. Page, and W. J. Baumol. Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences of the United States of America*, 101(46):16385–16389, 2004.

[4] M. LiCalzi and O. Surucu. The power of diversity over large solution spaces. *Manage. Sci.*, 58(7):1408–1421, July 2012.

[5] C. List and R. E. Goodin. Epistemic democracy: Generalizing the condorcet jury theorem. *Journal of Political Philosophy*, 9:277–306, 2001.

[6] L. S. Marcolino, A. X. Jiang, and M. Tambe. Diversity beats strength? - Towards forming a powerful team. In *15th International Workshop on Coordination, Organisations, Institutions and Norms (COIN 2013)*. (submitted).

[7] T. Obata, T. Sugiyama, K. Hoki, and T. Ito. Consultation algorithm for Computer Shogi: Move decisions by majority. In *Computers and Games'10*, pages 156–165, 2010.

[8] Pandanet. Introduction to Go. `http://www.pandanet.co.jp/English/introduction_of_go/`.