

Multi-Agent Team Formation: Diversity Beats Strength?

Leandro Soriano Marcolino, Albert Xin Jiang and Milind Tambe
University of Southern California, Los Angeles, CA, 90089, USA
{sorianom, 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? We propose a new model to address these questions. Our key contributions include: (i) we show that a diverse team can overcome a uniform team and we give the necessary conditions for it to happen; (ii) we present optimal voting rules for a diverse team; (iii) we perform synthetic experiments that demonstrate that both diversity and strength contribute to the performance of a team; (iv) we show experiments that demonstrate the usefulness of our model in one of the most difficult challenges for Artificial Intelligence: Computer Go.

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. Many works model team formation as selecting a team that accomplishes a certain task with the maximum expected value, given a model of the capabilities of each agent [Nair and Tambe, 2005; Guttman, 2008]. Other works go beyond a simple sum of skills, for example by considering synergetic effects in a team of agents [Liemhetcharat and Veloso, 2012] or studying how to automatically configure a network of agents [Gaston and desJardins, 2005].

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 [List and Goodin, 2001]. 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 [Obata *et al.*, 2011], deciding together stock purchases across different economic scenarios, or even picking items to recommend

to a large number of users [Burke, 2002]. 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 [Hong and Page, 2004; LiCalzi and Surucu, 2012]. However, [Hong and Page, 2004; LiCalzi and Surucu, 2012] 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. In the current literature on voting it is assumed a model where agents have a fixed probability to take the best action [Condorcet, 1785; List and Goodin, 2001; Young, 1995; Conitzer and Sandholm, 2005; Xia, 2011], and under that model it is not possible to show any advantage in having a diverse team of agents. Our experiments show, however, that a diverse team can outperform a uniform team of stronger agents. It is necessary to develop, therefore, a new model to analyze a team of voting agents.

In this work, we present a new model of diversity and strength for a team of voting agents. 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. Under this model, we can show that a team of diverse agents can perform better than a uniform team composed by strong agents. We present the necessary conditions for a diverse team to play better than a uniform team, and study optimal voting rules for a diverse team.

We present synthetic experiments with a large number of teams that demonstrate that both diversity and strength are important to the performance of a team. We also show results in one of the main challenges for Artificial Intelligence: Computer Go. Go is an iterative game, and the possible board states can represent a great variety of different situations, in such a way that the relative strength of different Go playing software changes according to the board state. Therefore, we can use our model to study a team of agents voting to play Computer Go. By using a diverse team we were able to increase the winning rate against Fuego (one of the strongest Go software) by 17%, and we could play 7% better than a team of copies of Fuego. Moreover, we could play 7% to 9% better than one of the versions of parallelized Fuego. We could also improve the performance of the diverse team by

12% using one of our proposed voting rules. Therefore, we effectively show that a team of diverse agents can have competitive strength, and even play better, than a uniform team composed by stronger agents. Our new model provides a theoretical explanation for our results.

2 Related Work

This work is related mainly to the study of team formation, diversity and voting. We will first introduce general works on team formation, then we will talk about diversity, and finally we will discuss voting.

Team formation is the problem of selecting the best possible team to accomplish a certain goal, given limited resources. In the traditional model, certain skills are necessary to accomplish a task, and we must select a team that has all the necessary skills with the maximum expected value [Nair and Tambe, 2005; Guttman, 2008]. More recent work go beyond a simple sum of skills and also models the synergy of a group [Liemhetcharat and Veloso, 2012], how to lead a group to the optimal joint action with a new ad-hoc agent [Agmon and Stone, 2012] or how to automatically configure a network of agents [Gaston and desJardins, 2005]. In [Matthews *et al.*, 2012], a team formation procedure is presented for a class of online soccer prediction games, and the system is able to play successfully against a large number of human players. However, the existing models do not cover the situation where we must select a team to vote together at each step of a complex problem. In this work, we present a new perspective to team formation, and we also introduce a new problem: in the pursuit of the best possible team of voting agents should we focus on the diversity of the team or on the strength of each individual member?

Hong and Page presented a contribution to team formation in the Social Science literature by showing the importance of diversity [Hong and Page, 2004]. They proposed a model for agents, and proved that a team of diverse agents can perform better than a team of high-ability agents. In their model, each agent has a set of local minima that they reach while trying to maximize an objective function. The agents can improve the solution from the local minima of their team members, therefore the search of a team stops only in the intersection of the local minima of all agents. By using a large number of diverse agents the system is able to converge to the optimal solution. Many papers followed their work [Luan *et al.*, 2012; Lakhani *et al.*, 2007; Krause *et al.*, 2011], showing the importance of diversity in different settings. Their model, however, does not cover situations where agents are unable to improve the solution from their team members' local minima. This can happen, for example, when we use existing software, that were not architected to collaborate in this way or when there are time constraints. Therefore, there are many situations where the agents have to collaborate in other ways, such as voting. If a team of agents votes, the system will not necessarily converge to an option in the intersection of their local minima. However, as we will show, it is still possible for a diverse team to play better than a uniform strong team.

A recent model to analyze diversity was proposed in [Licenzi and Surucu, 2012]. It is an equivalent model to Page's

and do not overcome the limitations previously described. In [Braouezec, 2010], the authors show the benefits of diverse agents voting to estimate the optimum of a single peaked function. In our work we are dealing with a harder problem, as the function to be optimized changes at every iteration. Another work that uses voting to study diversity is [West and Dellana, 2009], but they assumed that Page's model would work in a voting context, and do not propose a new model.

Concerning voting, the field has two possible views: voting for aggregating different preferences, and voting to estimate the best possible decision. Our work is related to the second view. The classical work in this line is the Condorcet's Jury Theorem [Condorcet, 1785]. According to the theorem, when facing a binary decision, as long as the average of the probability of each individual being correct is higher than $\frac{1}{2}$, a group of independent individuals doing majority voting will have a higher probability of being correct than the individuals alone. This theorem is extended to the k options case in [List and Goodin, 2001], where it is shown that if each of the individuals have a probability of choosing the best answer higher than choosing any other answer, the group performing majority voting will be stronger than the individuals alone. These theorems, however, do not present any benefits in having diverse agents. Researchers in Artificial Intelligence contributed to this view of voting by using a maximum likelihood approach to find the optimal voting rule. The idea is that given the votes, we can find which option has the highest probability of being the best, if we have a model of the probability distribution of the agents [Young, 1995; Conitzer and Sandholm, 2005; Xia, 2011]. However, they still do not address the issue of diversity and team formation in the context of voting, as they assume that all agents follow the same probability distribution. As all agents are essentially the same, team formation is not yet an issue in their work.

3 Methodology

Let Φ be a set of agents ϕ_i voting to decide an action a in the set of possible actions \mathbf{A} and Ω be a set of world states ω_j . We assume that we can rank the actions from best to worst and \mathbf{U}_j is the vector of expected utilities of the actions in world state ω_j , ordered by rank. The agents do not know the ranking of the actions, and will vote according to some decision procedure, characterized by a probability distribution function (pdf) over action ranks. Hence, each agent ϕ_i has a pdf $\mathbf{V}_{i,j}$ for deciding which action to vote for in state ω_j . Agents that have the same $\mathbf{V}_{i,j}$ in all world states will be referred as copies of the same agent.

Let α_j be the likelihood of world state ω_j . If the world states are equally frequent, we can use $\alpha_j = 1/|\Omega|$. We define strength as the weighted average of the expected utility of an agent or a team. It is given by the following dot product: $s = \sum_{\omega_j \in \Omega} \alpha_j \mathbf{V}_j \cdot \mathbf{U}_j$, where \mathbf{V}_j is the pdf of the agent/team in world state ω_j . \mathbf{V}_j can be calculated given a team of agents and a voting rule. A voting rule is a function that given the (single) votes of a team of agents, outputs an action.

We define the team formation problem as selecting from the space of all possible agents Ψ a set of n agents Φ that has the maximum strength in the set of world states Ω . An appli-

Agent	State 1	State 2	State 3	State 4	Strength
Agent 1	1	0	1	1	0.75
Agent 2	0	1	1	0	0.5
Agent 3	1	1	0	0	0.5
Agent 4	1	1	0	1	0.75
Agent 5	0	0	1	1	0.5

Table 1: A team of deterministic agents that can reach perfect play under simple voting. “1” indicates agent plays perfect action.

cation does not necessarily know $\mathbf{V}_{i,j}$ for all agents and for all world states. In this work, we will focus on showing that the naive solution of forming a team by selecting the strongest agents (or copies of the best agent) is not necessarily the optimal solution. Therefore, we are introducing a new problem to the study of team formation.

We define diversity as how different are the probability distributions of agents in Φ in the set of world states Ω : $d = \frac{1}{|\Phi|^2} \sum_{\omega_j \in \Omega} \sum_{\phi_i \in \Phi} \sum_{\phi_k \in \Phi} \alpha_j H(\mathbf{V}_{i,j}, \mathbf{V}_{k,j})$, where H is a distance measure between two pdfs. In this paper, we used the Hellinger Distance [Hellinger, 1909], given by:

$$H(\mathbf{V}_{i,j}, \mathbf{V}_{k,j}) = \frac{1}{\sqrt{2}} \sqrt{\sum_{a \in \mathbf{A}} (\sqrt{V_{i,j}(a)} - \sqrt{V_{k,j}(a)})^2}.$$

At each iteration, each agent will examine the current world state and submit its (single) opinion about which one should be the next action. The opinions are then combined using plurality voting, that picks as a winner the option that received the most votes. We consider in this paper three different voting rules: *simple* - break ties randomly, *static* - break ties in favor of the strongest agent overall, *optimal* - break ties in favor of the strongest agent of each world state. We consider the *static* voting rule because in some applications we might have a clear idea of which is the strongest agent overall, but the information of which is the strongest agent for a given world state might not be available. We will encounter this situation in the Computer Go domain, as will be clear in Section 4.2. This voting procedure will repeat at every iteration, until the end, when the system can obtain a reward.

3.1 Diversity Beats Strength

We present examples to demonstrate that a diverse team can play better than a uniform team. First, let’s consider the simplest case, when all agents are deterministic. The team made of copies of the strongest agent will play as well as the strongest agent, no matter how many members we add in the team. However, a team of diverse agents can overcome the strongest agent, and even reach perfect play, as we increase the number of agents. Consider, for example, the team in Table 1. This diverse team of 5 agents will reach perfect play under simple voting, while copies of the best agent (Agent 1 or Agent 4) will be able to play well only in 3 out of 4 world states, no matter how many agents we use in the team.

We can easily change the example to non-deterministic agents, by decreasing slightly the probability of them playing their deterministic action. A detailed description of the agents used in this example is available in the Appendix ¹.

¹In <http://teamcore.usc.edu/people/sorianom/ijcai13a.pdf>

The resulting strength of the teams is very similar to the deterministic case. Assuming all world states are equally likely, the strength of the diverse team is 0.9907, while copies of the best agent have strength 0.7499. Therefore, it is possible for a team of weak diverse agents to overcome a uniform team of stronger agents, when in certain states the individual agents are stronger than the overall strongest agent.

Even if we make the number of agents go to infinity, copies of the best agent will still be unable to perform the best action in one world state, and will play worse than the diverse team with only five agents. This situation is not considered in the Condorcet’s Jury Theorem, neither in the classical nor in the extended version, because they assume independent agents with a fixed pdf. Therefore, in the previous models, we would not be able to show the importance of diversity.

Necessary Conditions

We present a formal proof of the conditions necessary for a diverse team to play better than copies of the best agent, under the simple voting rule. If the conditions of the theorem are not met, we can simply use copies of the best agent as the optimal team. To simplify the presentation of the proof, we will consider a utility function with a value of 1 for the optimal action and 0 for the other actions. That is, we will consider the optimal team in a fixed world state as the team that has the highest probability of performing the optimal action. Let ψ_{best} be the strongest agent in Ψ , and a_{best} be the best action in a given world state.

Theorem 1. *For a diverse team to be the optimal team under the simple voting rule it is necessary that at least one agent in Ψ has a higher probability of taking the best action than ψ_{best} or a lower probability of taking a suboptimal action than ψ_{best} in at least one world state.*

Proof We develop the proof by showing that copies of the best agent of a given world state will be the optimal team in that world state. Therefore, it is necessary that the agents in the diverse team play better than the best agent overall in at least one world state. Let $\psi_{best,j}$ be the strongest agent in world state ω_j . Let’s define the pdf of this agent as $\langle p_1, \dots, p_k \rangle$, where p_1 is the probability of taking the best action. We will show that a team of n copies of $\psi_{best,j}$ doing simple voting will have a higher probability of taking the best action than a team of n agents composed of x copies of $\psi_{best,j}$ and m agents ψ_i doing simple voting, where the probabilities of each ψ_i are given by $\langle p_1 - \epsilon_i, p_2 + \gamma_{i2}, \dots, p_k + \gamma_{ik} \rangle$, $\gamma_{il} \geq 0 \forall l \in (2, k)$ and $\sum_{l=2}^k \gamma_{il} = \epsilon_i$.

Given a team of agents, let them all vote. We will start with a team of x copies of agent $\psi_{best,j}$. We will perform m iterations, and at each one we will add either another agent $\psi_{best,j}$ or agent ψ_i , where i is the current iteration. Let v_{i-1} be the current vote result. The result of v_{i-1} is either: (i) victory for a_{best} , (ii) tie between a_{best} and other options, (iii) defeat for a_{best} .

(i) If v_{i-1} is a victory for a_{best} , the new agent can change the result only when it votes for another option. Suppose a_l is an option that upon receiving one more vote will change a victory for a_{best} into a tie between a_{best} and a_l . Agent $\psi_{best,j}$ will vote for option a_l with probability p_l , while agent ψ_i will vote for option a_l with probability $p_l + \gamma_{il}$. Therefore, if v_{i-1}

is such that one vote can change a victory for a_{best} into a tie between a_{best} and other options, agent ψ_i will have a higher probability of changing a victory for a_{best} into a tie between a_{best} and other options.

(ii) If v_{i-1} is a tie between a_{best} and other options, $\psi_{best,j}$ will break the tie in favor of a_{best} with probability p_1 while ψ_i with probability $p_1 - \epsilon_i$. Therefore, agent $\psi_{best,j}$ will have a higher probability of breaking the tie in favor of a_{best} . Moreover, if a_l is an option that is currently tied with a_{best} , agent $\psi_{best,j}$ will vote for a_l with probability p_l , while agent ψ_i with probability $p_l + \gamma_{il}$. Therefore, agent ψ_i will have a higher probability of changing a tie between a_{best} and other options into a defeat for a_{best} .

(iii) If v_{i-1} is a defeat for a_{best} , $\psi_{best,j}$ will vote for a_{best} with probability p_1 while ψ_i will vote for a_{best} with probability $p_1 - \epsilon_i$. Therefore, if v_{i-1} is such that one vote can change a defeat for a_{best} into a tie between a_{best} and other options, agent $\psi_{best,j}$ will have a higher probability of changing a defeat for a_{best} into a tie between a_{best} and other options.

In all three cases, agent $\psi_{best,j}$ leads to a higher increase in the probability of picking a_{best} than agent ψ_i . Therefore, up to any iteration i , copies of $\psi_{best,j}$ will have a higher probability of playing the best action than a diverse team. Hence, if $\psi_{best,j} = \psi_{best} \forall j$, then copies of the best agent ψ_{best} will be the best team in all world states, and therefore it will be the optimal team. Therefore, for a diverse team to perform better, at least one agent must have either a higher probability of taking the best action or a lower probability of taking a suboptimal action than ψ_{best} in at least one world state. ■

This theorem, however, only gives the necessary conditions for a diverse team to be stronger than a non-diverse team. The sufficient conditions will depend on which specific game the agents are playing. Basically, given the pdf of the agents for a set of world states, we can calculate the pdf of both the diverse team, and the team made of copies of the best agent. If the diverse team has a higher probability of taking the best action in a subset of the world states that is enough for it to play better, considering that it will have a lower probability of taking the best action in the complementary subset, then the diverse team will play better than copies of the best agent.

Optimal Voting Rules

Given the result of a voting iteration, and the pdf of all agents in a world state, we can calculate which action has the highest probability of being the best. Formally, let $\langle p_1^i, \dots, p_k^i \rangle$ be the pdf of agent ϕ_i in world state ω_j , where p_l^i is the probability of playing the action with rank l . The optimal voting rule is the one that given a voting pattern, selects the action a_x that has the highest probability of being the best. Let \mathbf{Y} be the set of all other possible rank combinations for all other actions. We write an element of \mathbf{Y} as a sequence $y_1 \dots y_{k'-1}$, where y_l is a position in the ranking, and k' is the number of actions in the given voting pattern. We also define Φ_{a_1} as the set of agents in Φ that voted for an action a_1 , and \mathbf{A}_γ as the set of all actions in the given voting pattern. Assuming a uniform prior probability for the ranking of all actions, we can calculate the probability of a_x being the best action by:

$$\sum_{y_1 \dots y_{k'-1} \in \mathbf{Y}} \prod_{\phi_{i_x} \in \Phi_{a_x}} p_1^{i_x} \prod_{a_l \in \mathbf{A}_\gamma - \{a_x\}} \prod_{\phi_{i_l} \in \Phi_{a_1}} p_{y_l}^{i_l}$$

The derivation of the expression is available in the Appendix. We can extend this definition to picking the option with maximum expected utility by calculating the probability of each possible rank (instead of only the best one) and multiplying the resulting probability vector by a utility vector to obtain the expected utility of an option a_x .

However, it is possible to use a simpler voting rule. In our next theorem, we show that given some conditions, the optimal voting rule for a diverse team is to consider majority voting, but break ties in favor of the strongest agent that participates in the tie. Basically, we have to assume that all agents are strong enough to contribute to the team, so no agent should be ignored. If there are harmful agents in the team, we can try to remove them until the conditions of the theorem are satisfied. Again, we consider a utility function with a value of 1 for the optimal action and 0 for the other actions. Given a team Φ with size n , our conditions are:

Assumption 1 Weak agents do not harm

For any subset of Φ with an even number of agents n' , and for a fixed world state ω_j , let $\phi'_{best,j}$ be the best agent of the subset. We divide the agents in 2 sets: *Weak* containing the $n'/2 - 1$ agents that have the lowest probability of taking the best action and the highest probability of taking a suboptimal action, and *Strong* containing the $n'/2$ agents that have the highest probability of playing the best action and the lowest probability of taking a suboptimal action (except for the best agent $\phi'_{best,j}$, that is in neither one of the sets). We assume that when all agents in *Weak* and $\phi'_{best,j}$ vote together for an option a_x , and all agents in *Strong* vote together for another option a_y , the probability of a_x being the best action is higher than the probability of a_y being the best action.

Assumption 2 Strong agents are not overly strong

Given a fixed world state ω_j , we assume that if m_1 agents voted for an action a_x and m_2 agents voted for an action a_y , the probability of a_x being the best action is higher than a_y being the best action, if $m_1 > m_2$. If there is a situation where the opinion of a set of agents always dominates the opinion of another set, we can try to remove the dominated agents until the assumption holds true.

Theorem 2. *The optimal voting rule for a team is to consider the vote of all agents, but break ties in favor of the strongest agent if the above assumptions are satisfied.*

Proof Sketch Our detailed proof is available in the Appendix. We present here a proof sketch. By *Assumption 2* we know that we are looking for a tie-breaking rule, as the action chosen by the majority of the votes should always be taken. By *Assumption 1* we know that in the worst possible case, we should still break ties in favor of the strongest agent. If in the worst case, the group with the strongest agent still has a higher probability of selecting the best action than the group without the strongest agent, for any other case the group with the strongest agent will still have a higher probability. ■

As we are going to see in the next section, an application may not have the knowledge of the pdf of the agents in individual world states. Therefore, we also study an approximation of the optimal voting rule, that break ties in favor of the strongest agent overall, instead of breaking ties in favor of the strongest agent in a given world state. In the next section we

will see that both the optimal voting rule and our approximation improves the performance of a diverse team.

4 Results

4.1 Synthetic

We perform experiments using the quantal response (QR) model for the agents [McKelvey and Palfrey, 1995]. The quantal response model is a pdf from behavioral game theory to approximate how human beings (or non-rational players) behave while playing a game. It states that the probability of playing the best action is the highest, and it decays exponentially as the utility of the action gets worse. We use the QR model in our experiment, because it is a convenient way to represent non-rational agents with different strengths playing a game with a great number of options.

The pdf depends on a parameter, λ , that defines how rational (i.e., strong) is the agent. As λ gets higher, the agent provides a closer approximation to a perfect player. We define a λ_{ij} for each agent i and world state j .

We generated 1200 random teams of 4 agents, playing in 10 world states, and with 82 possible actions. We define each λ_{ij} as a random number in the interval $(0, 7)$, according to a uniform distribution. For each team, we can calculate the diversity and the average strength of the agents, according to the equations defined earlier. In Figure 1, we can see the performance of each team, as a function of diversity and the strength of its members. The strength of a team can be calculated after we generate the pdf of the team, by calculating the probability of all possible situations where the system would pick a particular ranking position. We assume that all world states are equally likely, hence the strength of a team is the average over all world states. We used a utility vector that gives a value close to 1 to the best action, and a low value to the other actions.

We performed a multiple linear regression for each voting rule. The following models were found: *simple*: $z = -0.09 + 1.48s + 0.45d$; *static*: $z = -0.03 + 1.36s + 0.55d$; *optimal*: $z = 0.09 + 0.92s + 1.29d$. The variable s is the average strength of the team members, d is the diversity of the team, and z is the strength of the team. The coefficient of multiple determination (R^2) of the models are 0.96, 0.81, 0.88, respectively.

As can be seen, both diversity and strength had a positive weight. This shows that groups with more diversity are stronger, given a fixed strength for their members. It is interesting to note that the impact of diversity increases as we change the voting rule from *simple* to *static*, and from *static* to *optimal*. The mean strength of all teams are $0.56(\pm 0.08)$, $0.61(\pm 0.08)$, $0.74(\pm 0.06)$, respectively. We can note that, as expected, *simple* had the lowest strength, followed by *static*, and *optimal* had the highest strength. The *optimal* voting rule is 30% stronger than *simple* voting in average.

4.2 Experiments in Computer Go

We perform experiments with four Go software: Fuego 1.1 [Enzenberger *et al.*, 2010], GnuGo 3.8, Pachi 9.01 [Baudiš and Gailly, 2011], MoGo 3 [Gelly *et al.*, 2006], and two (weaker) variants of Fuego (Fuego Δ and Fuego Θ), in a total

Player	# Higher P_{best}
GnuGo	17% (12%)
Pachi	21% (11%)
MoGo	20% (7%)
Fuego Δ	25% (6%)
Fuego Θ	26% (6%)

Table 2: Weak agents can play better in some board states. In parentheses, when the difference in P_{best} is 99% significant.

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 follow an UCT Monte Carlo Go algorithm [Gelly *et al.*, 2006]. The description of Fuego Δ and Fuego Θ is available at the Appendix. All results presented are obtained by playing 1000 9x9 Go games, in a HP dl165 with dual dodeca core, 2.33GHz processors and 48GB of RAM. Our system will play as white, against the original Fuego playing as black with opening database. Therefore, we will compare the winning rate of different agents and teams when playing against the same opponent. When we say that a result is significantly better than another, we use a *t-test* with 1% significance level ($\alpha = 0.01$).

Before introducing our results, we will use Theorem 1 to justify that it is worth it to explore a diverse team. If Fuego, the strongest agent, is always stronger in all board positions, then we can just use copies of Fuego as the optimal team. Therefore, we will first test if all agents are able to play better than Fuego in some board positions. In order to analyze the probability of an agent playing the best move in a certain board position, we need an estimate of the best move. Given a board position, we will ask Fuego to perform a movement, but we will give Fuego a time limit 50x higher than the default one. Therefore, Fuego is approximating how a perfect (or at least much stronger) player would play. We selected 100 board states, and we played all agents 50 times for each board state. Based on our estimate, we can calculate the probability of each agent playing the best move (P_{best}) for each board state. In Table 2, we can see in how many board states the agents have a higher P_{best} than Fuego (in its default time limit). As can be observed, all agents are able to play better than Fuego in some board positions, therefore it is possible for a diverse team to play better than copies of the best agent. As the number of board states where an agent plays better is not small, we can expect that a diverse team should be able to overcome the uniform team.

According to Theorem 2, if we assume that the weak agents (like GnuGo) are not weak enough to harm the system, and the strong agents (like Fuego and its variants) are not strong enough to dominate a subset of the agents, then the optimal voting rule is to break ties in favor of the strongest agent. However, during a game the system does not have access to the pdf of the agents, and has no way to identify which is the strongest agent. Therefore, we present results using the *static* voting rule, that break ties in favor of the strongest agent overall. Based on our synthetic results, we can predict that *static* should perform better than *simple*. We also tried a weighted voting rule, which allowed us to empirically learn the best

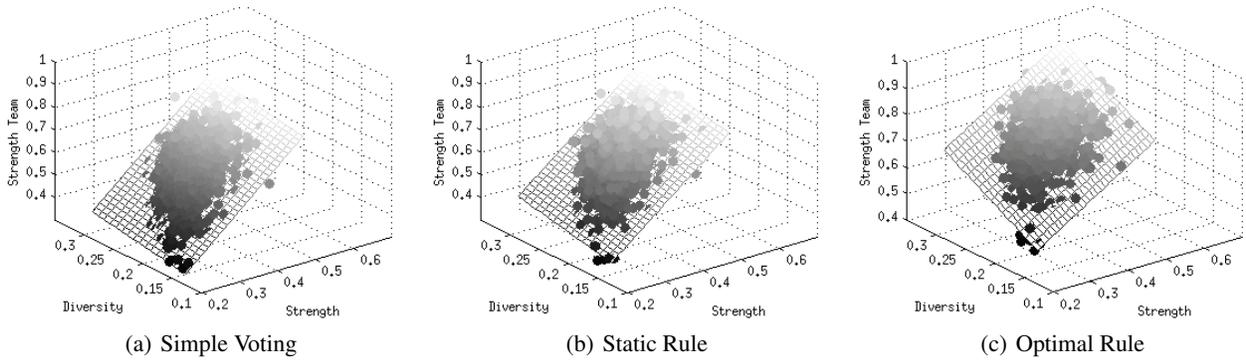


Figure 1: 1200 random teams of 4 agents.

weights by a hill climbing algorithm. The resulting rule was equivalent to the *static* voting rule.

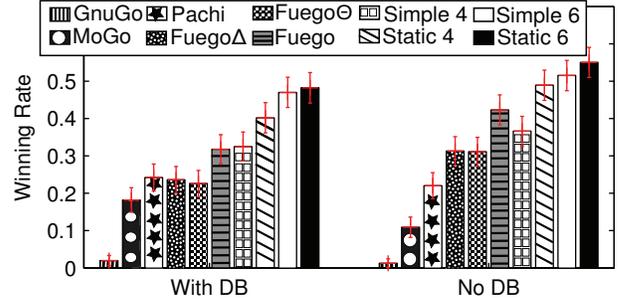
We call a team composed by different Go software as “Diverse” or by the name of the voting rule that they use (“Simple” or “Static”). The team of copies of the strongest agent (Fuego) will be called “Copies”. The copies are initialized with different random seeds, therefore due to the nature of the search algorithms, they will not always choose the same movement. When we want to be explicit about the number of agents in a team we will add a number after the name of the team. “Diverse” is composed by Fuego, GnuGo, Pachi and MoGo when executed with four agents, and is composed by all agents when executed with six agents. We also work with a parallelized version of Fuego (“Parallel”), and we will add a number after its name to indicate the number of threads.

We can see a summary of all our results in Figure 2. Diverse plays significantly better than Fuego, with 6 agents or when using the *static* voting rule. When we keep the opening database, Diverse plays significantly better than Copies and Parallel with 6 agents. Without the opening database, Diverse still plays significantly better than Parallel with 6 agents, but the difference between Diverse and Copies is not statistically significant. Static is either significantly better than Simple, or the difference between them is not statistically significant.

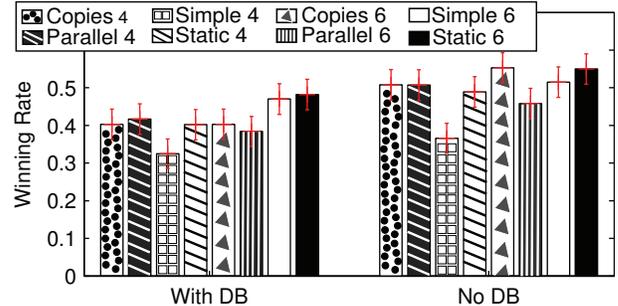
By the classical view of voting, our experimental result is not expected. If we view each agent as having a fixed pdf, we would predict that copies of the best agent would perform much better than a diverse team with weaker agents. However, in our results we showed that the diverse team has a competitive strength, and is able to play even better than copies of the best agent when we keep the opening database. Our model provides a theoretical explanation for our results.

5 Conclusion

This paper opens a new path for multi-agent team formation. Unlike existing work, we study team formation in the context of agents that vote together at each step of a complex problem. This introduces a new conflict between focusing on the diversity of a team or on the strength of each individual member when we pursue the best possible team. We opened this new path by showing that the naïve solution of selecting the strongest agents does not work, as it is possible for a team of



(a) Winning rate of single agents and the Diverse Team (Simple/Static)



(b) Winning rate of the uniform team (Copies), the Diverse team (Simple/Static), and a parallelized agent (Parallel)

Figure 2: Results in the Computer Go domain. The error bars show the confidence interval, with 99% of significance.

weak but diverse agents to perform better than a uniform team made of copies of the strongest agent. We also introduced the study of optimal voting rules for a diverse team, and showed both via mathematical analysis and experimentation that our proposed rules improve the performance of a team. Our experiments with a set of Go playing agents demonstrate the importance of our model, as the model provides a theoretical explanation of our results that would not be possible otherwise. It is also a concrete example that diversity can overcome strength in multi-agent team formation.

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

References

- [Agmon and Stone, 2012] Noa Agmon and Peter Stone. Leading ad hoc agents in joint action settings with multiple teammates. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '12*, pages 341–348, Richland, SC, 2012.
- [Baudiš and Gailly, 2011] Petr Baudiš and Jean-loup Gailly. Pachi: State of the Art Open Source Go Program. In *Advances in Computer Games 13*, November 2011.
- [Braouezec, 2010] Yann Braouezec. Committee, expert advice, and the weighted majority algorithm: An application to the pricing decision of a monopolist. *Computational Economics*, 35(3):245–267, March 2010.
- [Burke, 2002] Robin Burke. Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4):331–370, November 2002.
- [Condorcet, 1785] Marquis de Condorcet. *Essai sur l'application de l'analyse a la probabilité des décisions rendues a la pluralité des voix*. L'Imprimerie Royale, 1785.
- [Conitzer and Sandholm, 2005] Vincent Conitzer and Thomas Sandholm. Common voting rules as maximum likelihood estimators. In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence, UAI'05*, pages 145–152. Morgan Kaufmann Publishers, 2005.
- [Enzenberger et al., 2010] M. Enzenberger, M. Müller, B. Arneson, and R. Segal. Fuego - An open-source framework for board games and go engine based on Monte Carlo Tree Search. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(4):259–270, dec. 2010.
- [Gaston and desJardins, 2005] Matthew E. Gaston and Marie desJardins. Agent-organized networks for dynamic team formation. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '05*, pages 230–237, New York, NY, USA, 2005. ACM.
- [Gelly et al., 2006] Sylvain Gelly, Yizao Wang, Rémi Munos, and Olivier Teytaud. Modification of UCT with patterns in Monte-Carlo Go. Technical report, 2006.
- [Guttman, 2008] Christian Guttman. Making allocations collectively: Iterative group decision making under uncertainty. In Ralph Bergmann, Gabriela Lindemann, Stefan Kirn, and Michal Pechoucek, editors, *Proceedings of the 6th German Conference on Multiagent System Technologies*, volume 5244 of *Lecture Notes in Computer Science*, pages 73–85, Kaiserslautern, Germany, 2008. Springer.
- [Hellinger, 1909] Ernst Hellinger. Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. *Journal für die reine und angewandte Mathematik*, 136:210–271, 1909.
- [Hong and Page, 2004] Lu Hong and Scott E. Page. 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.
- [Krause et al., 2011] Stefan Krause, Richard James, Jolyon J. Faria, Graeme D. Ruxton, and Jens Krause. Swarm intelligence in humans: diversity can trump ability. *Animal Behaviour*, 81(5):941–948, May 2011.
- [Lakhani et al., 2007] Karim R. Lakhani, Lars Bo Jeppesen, Peter A. Lohse, and Jill A. Panetta. The value of openness in scientific problem solving. *HBS Working Paper*, (07-050), 2007.
- [LiCalzi and Surucu, 2012] Marco LiCalzi and Oktay Surucu. The power of diversity over large solution spaces. *Management Science*, 58(7):1408–1421, July 2012.
- [Liemhetcharat and Veloso, 2012] Somchaya Liemhetcharat and Manuela Veloso. Modeling and learning synergy for team formation with heterogeneous agents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '12*, pages 365–374, Richland, SC, 2012.
- [List and Goodin, 2001] Christian List and Robert E. Goodin. Epistemic democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy*, 9:277–306, 2001.
- [Luan et al., 2012] Shenghua Luan, Konstantinos V.atsikopoulos, and Torsten Reimer. When does diversity trump ability (and vice versa) in group decision making? A simulation study. *PLoS One*, 7(2):e31043, 2012.
- [Matthews et al., 2012] Tim Matthews, Sarvapali D. Ramchurn, and Georgios Chalkiadakis. Competing with humans at Fantasy Football: Team formation in large partially-observable domains. In *Proceedings of the 26th Conference of the Associations for the Advancement for Artificial Intelligence*, pages 1394–1400, 2012.
- [McKelvey and Palfrey, 1995] Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. In *Games and Economic Behavior*, volume 10, pages 6–38. Elsevier, 1995.
- [Nair and Tambe, 2005] Ranjit Nair and Milind Tambe. Hybrid BDI-POMDP framework for multiagent teaming. *Journal of Artificial Intelligence Research*, 23(1):367–420, April 2005.
- [Obata et al., 2011] Takuya Obata, Takuya Sugiyama, Kuni-hito Hoki, and Takeshi Ito. Consultation algorithm for Computer Shogi: Move decisions by majority. In *Computer and Games'10*, volume 6515 of *Lecture Notes in Computer Science*, pages 156–165. Springer, 2011.
- [West and Dellana, 2009] David West and Scott Dellana. Diversity of ability and cognitive style for group decision processes. *Information Sciences*, 179(5):542–558, 2009.
- [Xia, 2011] Lirong Xia. *Computational voting theory: game-theoretic and combinatorial aspects*. PhD thesis, Durham, NC, USA, 2011. AAI3469020.
- [Young, 1995] Peyton Young. Optimal voting rules. *Journal of Economic Perspectives*, 9(1):51–64, 1995.