

A Detailed Analysis of a Multi-agent Diverse Team

Leandro Soriano Marcolino, Chao Zhang, Albert Xin Jiang, and Milind Tambe

University of Southern California, Los Angeles, CA, 90089, USA
{sorianom, zhan661, jiangx, tambe}@usc.edu

Abstract. In an open system we can have many different kinds of agents. However, it is a challenge to decide which agents to pick when forming multi-agent teams. 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 Go¹.

Keywords: Multi-agent Systems, Coordination & Collaboration, Distributed AI

1 Introduction

It is well known that teams can often outperform individual agents. However, different combinations of agents have different performances, and it is even possible for a team to perform worse than its members. In an open system, we can have a variety of agents available. How can we pick a limited number of them in order to form strong teams?

After forming a team, their members must work together. Voting is an important coordination mechanism for a multi-agent team. By voting, we can use agents that were not originally designed to work together, and there are theoretical guarantees that a team of voting agents can get closer to finding the best possible decision in a given situation [2]. Sometimes the agents must vote continuously in many different scenarios. Consider, for example, agents that are cooperating in a board game [3], deciding together stock purchases across different economic scenarios, or even picking items to recommend to a large number of users [4]. This situation imposes a conflict for team

¹ This paper is an extended version of [1]. We include here more empirical results: while in [1] there are results only for white, here we also study in Section 4.2 teams playing as black, hence showing a more general result. Moreover, in [1] we analyze the agents only using our proposed model, but here (again in Section 4.2) we also analyze them using classical voting models, emphasizing the importance of our new model. In addition, we present in Section 5 a new study by a human expert of some games from our experiments, in order to better understand why a team of diverse agents is able to overcome a uniform team of strong agents. Finally, we have new discussions in the Conclusion.

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 [5, 6]. However, [5, 6] 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 [7, 2, 8–10], 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.

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 show 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 18.7%, and we could play 11% better than a team of copies of Fuego. Moreover, we could play 15.8% better than one of the versions of parallelized Fuego. We could also improve the performance of the diverse team by 12.7% 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 minimum cost [11, 12]. More recent work goes beyond a simple sum of skills and also models the synergy of a group [13], how to lead a group to the optimal joint action with a new ad-hoc agent [14] or how to automatically configure a network of agents [15]. In [16], a team formation procedure is presented for a class of online football 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 [5]. 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 [17–19], 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 more recent model to analyze diversity was proposed in [6]. It is an equivalent model to Page’s and still do not overcome the limitations previously described. In [20], 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 [21], 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 [7]. 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 plurality voting will have a higher probability of being correct than the individuals alone. This theorem is extended to the k options case in [2], 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 plurality 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 [8–10]. 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 we expect the world states to be 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 application 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 naïve 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 [22], 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.

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.

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.

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 (in <http://teamcore.usc.edu/people/sorianom/coin2013Book-appendix.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.

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} in at least one world state, 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 in 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 in option a_l with probability p_l , while agent ψ_i will vote in 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, agent $\psi_{best,j}$ will break the tie in favor of a_{best} with probability p_1 while agent 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} , agent $\psi_{best,j}$ will vote for a_{best} with probability p_1 while agent 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.

Now we study 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_l} as the set of agents in Φ that voted for an action a_l , 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:

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

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 plurality 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 in an option a_x , and all agents in *Strong* vote together in 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 in an action a_x and m_2 agents voted in 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 most 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. ■

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 synthetic experiments using the quantal response (QR) model for the agents [23]. 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.

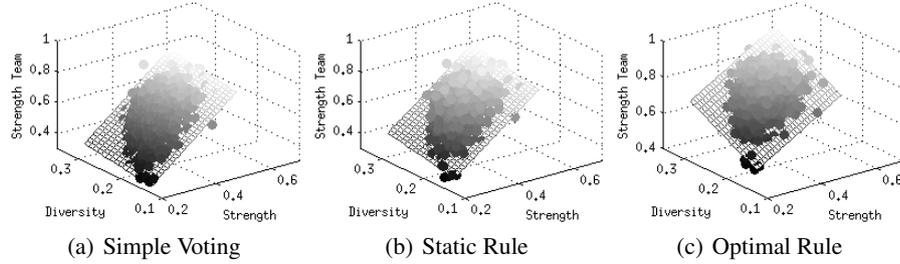


Fig. 1. 1200 random teams of 4 agents.

4.2 Experiments in Computer Go

We also perform experiments with four Go software: Fuego 1.1, GnuGo 3.8, Pachi 9.01, MoGo 3, and two (weaker) variants of Fuego (Fuego Δ and Fuego Θ), 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 follow a UCT Monte Carlo Go algorithm [24]. 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. We first present results when our system plays as white, against the original Fuego playing as black with opening database. Then, we present results of our system playing as black, against the original Fuego playing as white with opening database. 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 ($\alpha = 0.01$).

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 4 agents, and is composed by all agents when executed with 6 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.

Before introducing our results, we first analyze the agents under the classical voting theory and under our proposed theory. To simplify the analysis, we consider here the probability of playing the best move (P_{best}); therefore, we consider a utility vector with

a value of 1 for the best move, and 0 for the other moves. We start by the classical voting theories. In order to estimate P_{best} , we use 1000 board states from our experiments. In 1000 games, we randomly choose a board state between the first and the last movement. We then ask Fuego to perform a movement in that state, but we 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. To avoid confusion with the names we will call this agent Perfect. We then obtain Perfect’s evaluation for all the positions of the board, and organize them into a ranking.

We ran all agents in the selected 1000 board states and for each state we verify in which position of the ranking each agent would play. If, instead of playing, the agent resigns, we randomly pick a different board state and regenerate the data for all agents, including Perfect’s evaluation. Based on that, we can generate a histogram for all agents. Some examples can be seen in Figure 2. Although we do not have enough space to show here all histograms, they have similar shapes, giving a high frequency for the best possible move and lower frequencies for the subsequent moves, with the stronger agents having a higher frequency of playing the best move. The interested reader can see all the generated histograms in the Appendix.

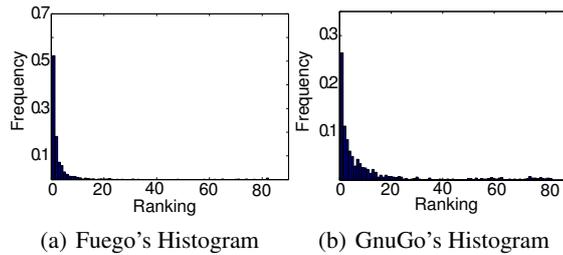


Fig. 2. Histogram of the agents, using real data.

Assuming that the agents are independent, and that each one will choose a move according to the probability distribution corresponding to its histogram, we can calculate P_{best} of any group and voting rule that we want. Basically we have to calculate the probability of all the possible situations where the system would pick the best move. For a team of k agents we have to calculate $O(n^{k-1})$ probabilities, where n is the number of possible options. While for a team of 4 agents we are able to calculate the precise value, for a team of 6 agents we are going to show approximations.

In Table 2 we can see P_{best} of each individual player and of all teams. The P_{best} of the teams is higher than the P_{best} of each one of the agents, and is higher for a team of 6 agents than for a team of 4 agents. This result is expected when we consider the extended version of the Condorcet’s Jury Theorem [2], at least for a uniform team. According to the theorem P_{best} approaches 1 when the number of agents goes to infinity. However, we would also expect Copies to perform better than Diverse. Would it be possible, then, for a diverse team to perform better than a uniform team?

(a) Players		(b) Teams	
Player	P_{best}	Team	P_{best}
Fuego	52.3%	Simple 4	57.5%
GnuGo	26.4%	Static 4	61.8%
Pachi	40.6%	Copies 4	79.6%
MoGo	40.8%	Simple 6	71.1%
Fuego Δ	48.8%	Static 6	72.4%
Fuego Θ	47.7%	Copies 6	86.6%

Table 2. Probability to select the best move of each player and the teams.

Intuitively, we would expect that a uniform team would agree on certain moves much more often than a diverse team. And indeed, when we look at the graph of the frequency of the size of the set of agents that voted for the winning move (Figure 3), we can see that they are very different. In the x-axis we show the number of agents that agreed in the selected movement, and in the y-axis the frequency of each number considering all moves in the 1000 games. The expected size of the set for Diverse is 3.50, while for Copies is 4.43. Therefore, if Fuego plays badly in a certain board state, all copies of Fuego would also tend to vote for the same bad moves. In a diverse team, however, some agents could be able to play better in that particular situation. The extended Condorcet’s Jury Theorem assumes that agents are independent, but in fact their relative performances might change according to the state of the board.

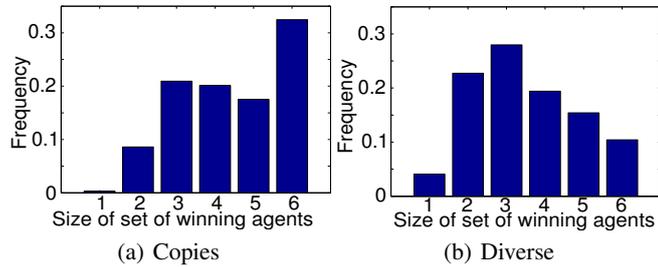


Fig. 3. Expected size of the set of agents that vote for the winning move, with 6 agents and no opening database.

Now we analyze the agents according to our proposed theory. 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 test if all agents are able to play better than Fuego in some board positions. We selected 100 board states, and we played all agents 50 times for each board state. Based on our estimate of the best move (obtained from Perfect), we can calculate P_{best} for each agent and for each board state. In Table 3, 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.

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

Table 3. Weak agents can play better in some board states. In parentheses, we show when the difference in P_{best} is 99% significant.

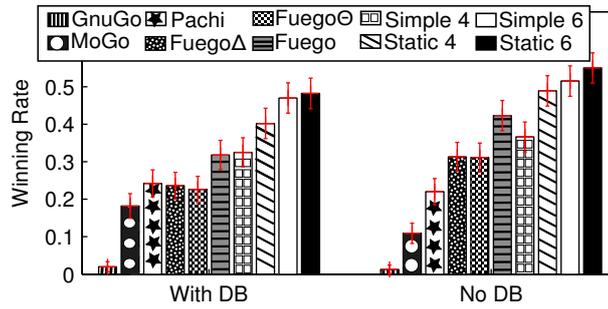
According to Theorem 2, if we assume that the weak agents (like GnuGo) are not weak enough to harm the system, and the strong agents (like Fuego and its variants) are not strong enough to dominate a subset of the agents, then the optimal voting rule is to break ties in favor of the strongest agent. However, 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 weights by a hill climbing algorithm. The resulting rule was equivalent to the *static* voting rule.

We can see our results for white in Figure 4(a,b). Diverse plays significantly better than Fuego, with 6 agents or with 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 significant. Static is either significantly better than Simple, or the difference between them is not significant.

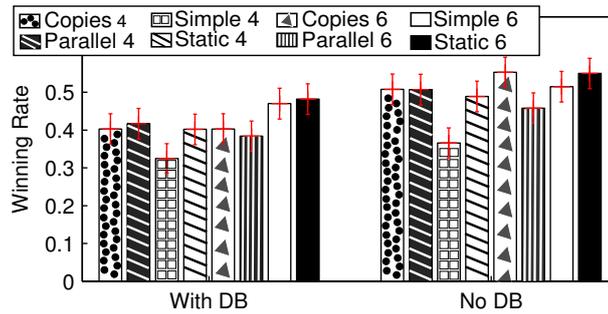
In Figure 4(c,d) we can see the results for black. Again, Diverse plays significantly better than Fuego when using the *static* voting rule. This time, however, Diverse (with 6 agents or using the *static* voting rule) is able to play significantly better than Copies without the opening database, but with the opening database the difference between them is not significant. Again, Static is either significantly better than Simple, or the difference between them is not significant. Static is always significantly better than Parallel.

To verify the generality of improving the results by the *static* voting rule and by adding more agents, we also played our system as white against Pachi as black, without opening database. Simple 4 won 56.2% of the games, Static 4 won 65.5% and Simple 6 won 66.8%. Therefore, these techniques can improve the results in other situations.

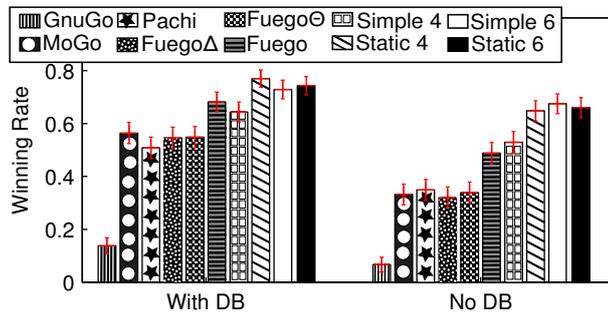
By the classical view of voting, our experimental result is not expected. If we view each agent as having a fixed pdf, we would predict that copies of the best agent would perform much better than a diverse team with weaker agents. However, in our results we



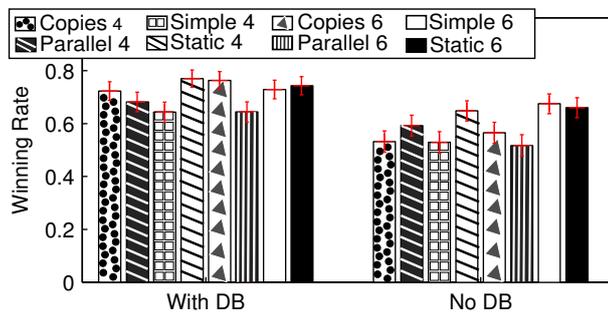
(a) Results for white. Single agents and the diverse team (Simple/Static).



(b) Results for white. The uniform team (Copies), the diverse team (Simple/Static), and a parallelized agent (Parallel).



(c) Results for black. Single agents and the diverse team (Simple/Static).



(d) Results for black. The uniform team (Copies), the diverse team (Simple/Static), and a parallelized agent (Parallel).

Fig. 4. Results in the Computer Go domain. The error bars show the confidence interval, with 99% of significance.

showed that the diverse team has a competitive strength, and is able to play even better than copies of the best agent in some situations. Our new model provides a theoretical explanation for our experimental results.

5 Detailed Study: Why a Team of Diverse Agents Perform Better?

We study in detail three games from our experiments, in order to better understand why a team of weak players can perform as well, or better than a team made of copies of the best player. We study games with 6 agents, using the *simple* voting rule. According to our theoretical work, at least one agent must play better than the strongest agent in at least one world state for a diverse team to overcome a uniform team. These are only necessary conditions, for a diverse team to effectively play better this must happen in many world states, specially in critical situations that can decide the game. Here we show that this really happens in Computer Go, based on an analysis by an expert human player. As Go is a complex game, we note that some expert readers might not agree completely with all points of our analysis. Although we present results in the Computer Go domain, this phenomenon should also occur in other complex domains, where the relative strength of the agents change according to the world state.

These games are analyzed by Chao Zhang, a 4-dan amateur Go player. In order to show that the weak agents are not playing better simply by chance, we estimate the probability of all agents playing all analyzed moves by repeatedly playing them 100 times in the board state under consideration. Based on these probabilities, we calculate the probabilities of the diverse team and the uniform team, to show that the diverse team would perform better in these board states. An important point to note is that it is not the case that a certain subset of the agents always vote for a better move; the set of agents that can find a better move than Fuego changes according to each board state.

This analysis requires some Go knowledge to be fully understood. 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 (i.e. they are "killed"). 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 [25].

We first analyze the Go game in Figure 5. In some positions, the weak agents vote for better moves than Fuego, the strongest agent. Move 11 is a very interesting situation. Here, Fuego, Pachi and MoGo vote for move D4, while GnuGo votes for E8 (X). Even though GnuGo is the weakest agent, in this situation it is able to find a better move than all other agents. E8 is better because it allows white to get the territory in the upper left corner. Besides, white can aim at G7 to kill the black group in the upper right. If white plays D4, black can play E8 to kill white aiming at the upper left corner. Unfortunately, GnuGo loses the vote in this situation. In all other positions, we show situations where the weak agents vote together for a better move than Fuego. For example, in move 23, Fuego votes for B4 (Δ) while Pachi, MoGo, Fuego Δ and Fuego Θ vote for B7. If white chooses B7, white can kill C7&D7 or B5&C5&D5&E5. If black saves C7&D7, white

can use B4 to kill the other group; If black saves B5&C5&D5&E5, white can use C8 to kill C7&D7. If white chooses B4, black will use B7 to kill the white group in the upper left. Fuego's mistake is critical in this situation, and would lead to losing the game. In move 45, Fuego would make another mistake. Fuego votes for B9 (Δ), while GnuGo and Pachi vote for H3. B9 wastes a move: it cannot affect the final result and wastes a chance for further developments. H3, on the other hand, is aiming at killing black in the right bottom. In this case Fuego Δ also votes for B9, so the *static* voting rule would choose a worse move. In move 63, Fuego would play E3, while GnuGo, Pachi, Fuego Δ and Fuego Θ vote for G2. If white plays G2, the black group in the right bottom dies, while if white plays E3, white cannot kill them. This is another critical mistake, that would make white lose the game. Finally, in move 75, Fuego votes for A7, while Pachi and Fuego Θ vote for G2. G2 is better than A7, as it allows white to have a larger territory. As can be seen, there are many situations where the weaker agents vote together for a better move than Fuego. The probabilities of each agent playing the analyzed moves can be seen in Table 4. It is clear that Fuego did not choose the worse move by accident: in many cases it has a lower probability than the other agents of playing the best move between the two options. Consequently, the uniform team is still not able to perform well in these situations, it still has a low probability of playing the best move, and it is always outperformed by the diverse team. In some situations, the probability of playing the worst move even increases by using multiple copies of Fuego.

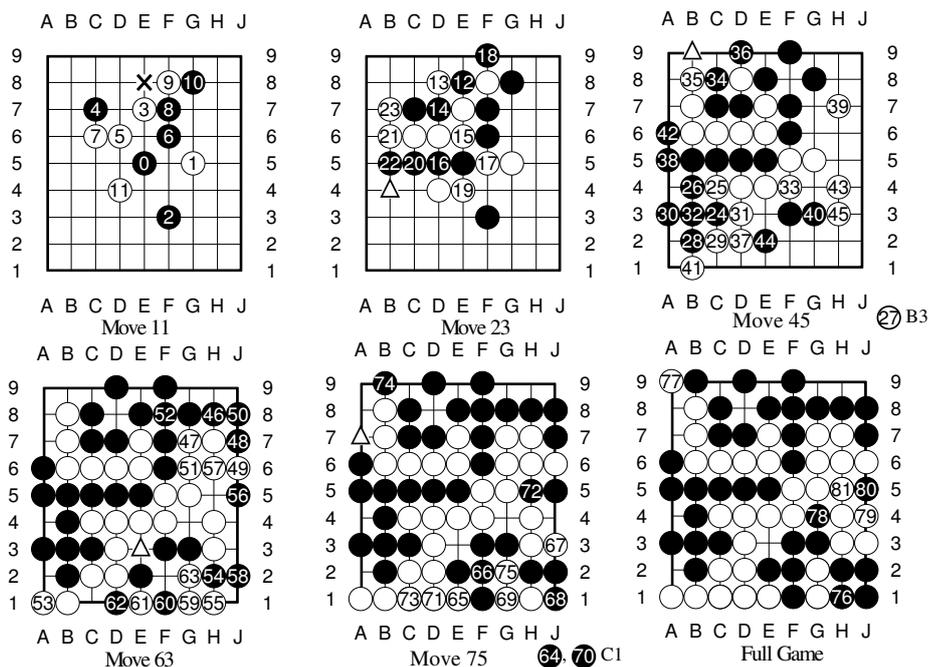


Fig. 5. First example, the diverse team plays as white without the opening database against Fuego. White wins by resignation.

Agent	Move 11		Move 23		Move 45		Move 63		Move 75	
	E8*	D4	B7*	B4	H3*	B9	G2*	E3	G2*	A7
Fuego	2%	51%	83%	14%	1%	76%	53%	29%	22%	16%
GnuGo	100%	0%	0%	0%	100%	0%	100%	0%	0%	0%
Pachi	6%	75%	30%	70%	46%	0%	78%	1%	35%	1%
MoGo	2%	61%	100%	0%	0%	0%	0%	84%	53%	0%
Fuego Δ	24%	19%	100%	0%	16%	19%	76%	13%	24%	7%
Fuego Θ	35%	9%	99%	0%	12%	30%	78%	10%	31%	11%
Diverse	15%	57%	99%	0%	20%	28%	88%	7%	45%	5%
Copies	0%	73%	95%	4%	0%	98%	63%	26%	23%	21%

Table 4. Probability of playing the moves in the first example. * indicates the better move.

We now analyze the game in Figure 6. In move number 4, Fuego Δ and Pachi vote for C7, while Fuego votes for move G3 (Δ). G3 is a bad opening for white, because the two white groups would be split by black. Another example is in move 7, when GnuGo, Pachi and Fuego Δ vote for B6, while Fuego votes for G7 (Δ). Black and white are fighting in the upper left corner. If white plays G7, it waives the fight and plays in a place that is not immediately important. White should choose B6 to continue the fight in order to win. Even GnuGo, the weakest agent, knows that B6 is a better move. In move 25, GnuGo and Mogo choose A8, while Fuego chooses F2 (Δ). If white does not play A8, black will play A5 to kill the white group in the left side. White has to kill with A8. This time Fuego's mistake is critical, and could lead to losing the whole game. In this situation GnuGo helps avoid a critical mistake, because Fuego Θ also votes for F2. Moreover, it is an example of a case where the *static* voting rule fails, as it would break the tie in favor of Fuego. We expect that significant improvements in game play would be possible if we learn which is the strongest agent in a given situation, and better approximate the *optimal* voting rule. Another interesting move is 37. Fuego and Fuego Δ vote for D2, while MoGo and Fuego Θ vote for E3. Both moves are equally good, as they get the same territory. However, GnuGo might have a better move: F6. If white plays F6, it can aim at both G6 and F4 for the next moves, which will cause great harm to black's territory. This is another example of a situation where the weakest agent has a better move than all other agents. The probabilities of each agent playing the analyzed moves can be seen in Table 5. Again, we can see that the diverse team would have a higher probability of finding the better moves than the uniform team.

In the games with the opening database, an interesting one is in Figure 7. In move 29, GnuGo, Pachi and MoGo choose D2, while Fuego votes for B8 (Δ). D2 can protect the lower left, while B8 cannot kill the black group in the upper left, and ends up making it more solid. In move 31, Pachi and MoGo vote for B3, and Fuego votes for B1 (Δ). Even if both moves might be able to kill the black stone in B2, B3 can kill it for sure. If white plays B1, black can play B3 and would lead to complications. This mistake could make white lose the game. If black survives, it can kill the white group in the lower left. In move 45, Pachi, MoGo and Fuego Δ vote for F4, while Fuego votes for A4 (Δ). F4 splits black into two groups and can make use of this division in the future. A4 just wastes a move and gives black more territory. In move 51, Pachi and MoGo choose E9, and Fuego chooses H6 (Δ). E9 makes the white group on the left survive, while H6

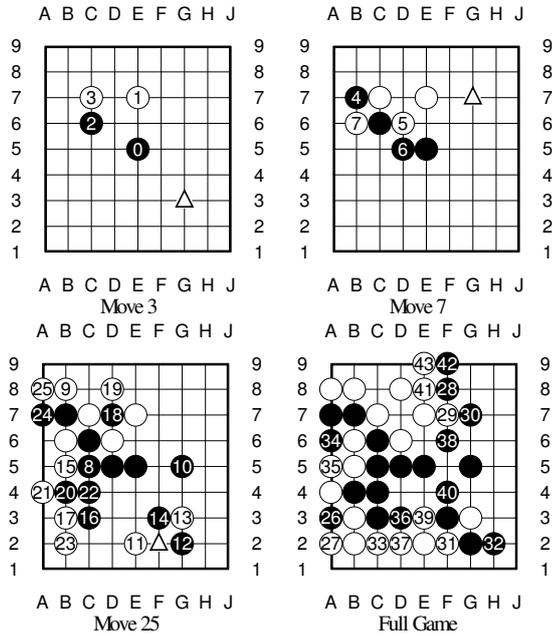


Fig. 6. Second example, the diverse team plays as white without the opening database against Fuego. White wins by resignation.

Agent	Move 3		Move 7		Move 25		Move 37		
	C7*	G3	B6*	G7	A8*	F2	F6*	D2	E3
Fuego	20%	2%	7%	41%	11%	30%	1%	53%	19%
GnuGo	0%	0%	100%	0%	100%	0%	100%	0%	0%
Pachi	27%	14%	99%	1%	28%	19%	26%	27%	0%
MoGo	0%	8%	1%	0%	89%	0%	0%	41%	45%
Fuego Δ	20%	0%	34%	20%	28%	10%	0%	83%	7%
Fuego \ominus	25%	4%	50%	7%	37%	11%	0%	80%	12%
Diverse	-%	-%	77%	0%	70%	1%	1%	90%	5%
Copies	19%	0%	2%	30%	7%	27%	0%	84%	8%

Table 5. Probability of playing the moves in the second example. Some results are unavailable due to lack of memory. * indicates the better move.

wastes a move and will lead to the death of the white group. This is a critical mistake, that would make white lose the game. In move 67, GnuGo and MoGo vote for B4, while Fuego votes for G2 (Δ). B4 is better, as it can get more territory. G2 just wastes a move. The probabilities of each agent playing the analyzed moves can be seen in Table 6. Again, in all these situations the diverse team has a higher probability of playing the better move than the uniform team. In some cases, the probability of playing the worse move even increases with multiple copies of Fuego.

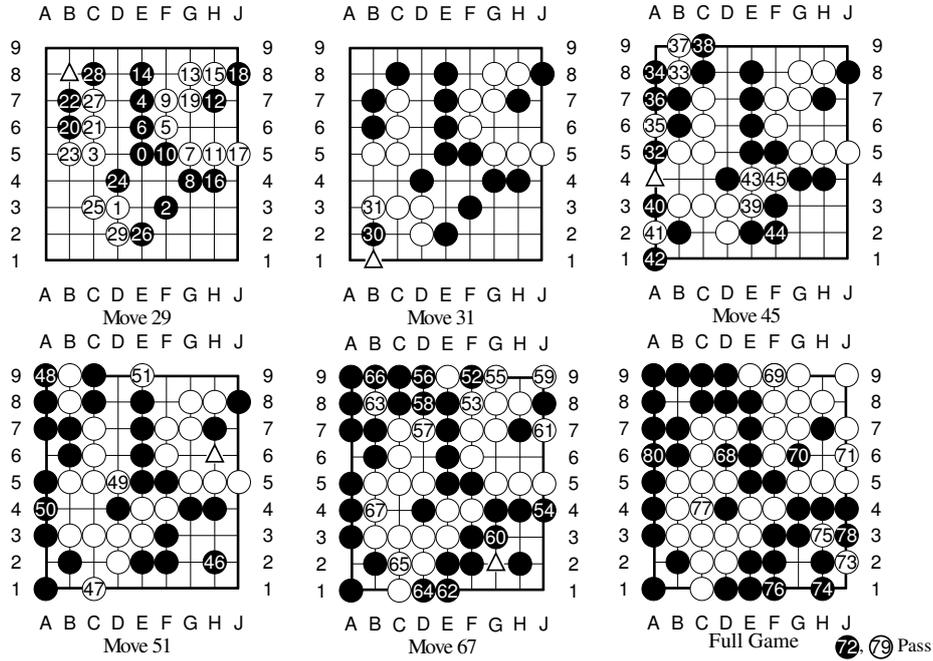


Fig. 7. Third example, the diverse team plays as white with the opening database against Fuego. White wins by resignation.

Agent	Move 29		Move 31		Move 45		Move 51		Move 67	
	D2*	B8	B3*	B1	F4*	A4	E9*	H6	B4*	G2
Fuego	3%	16%	44%	26%	17%	40%	0%	35%	0%	12%
GnuGo	100%	0%	0%	0%	0%	0%	0%	0%	100%	0%
Pachi	77%	18%	64%	0%	78%	17%	90%	0%	3%	0%
MoGo	91%	0%	98%	0%	92%	0%	51%	0%	46%	4%
Fuego Δ	6%	4%	11%	1%	51%	12%	0%	1%	0%	9%
Fuego Θ	5%	7%	13%	0%	50%	21%	0%	2%	0%	5%
Diverse	82%	3%	75%	0%	54%	12%	37%	0%	9%	1%
Copies	0%	12%	56%	32%	5%	53%	0%	44%	0%	4%

Table 6. Probability of playing the moves in the third example. * indicates the better move.

6 Conclusion and Discussions

We showed that diverse teams can outperform teams composed by copies of the best player. However, it is still a challenge to find the best possible teams. In an open multi-agent system the pdfs of the agents are generally not available. Moreover, in many complex scenarios we cannot even easily enumerate all the possible states of the world. Hence, given a world state, how can we quickly and automatically know the relative strength of the different agents? This is still an important open problem.

We gave an initial step by studying in detail different scenarios where diverse agents are able to outperform the best agent. One possible direction for future work is to identify common characteristics of world states where a certain agent is able to play better than the best agent. Hence, given a new world state we would be able to estimate the strongest agent for that specific world state and better approximate the *optimal* voting rule. In addition, we could also dynamically change the team in order to have the best (or close to the best) possible one for each different scenario.

In real-life scenarios, like robot teams, the problem is even more challenging. We can always estimate the pdf of an agent by running it multiple times in a given world state, if we have at least an estimation of the ground truth. However, for an embodied agent, the number of times we can sample might be very limited. A similar challenge is faced in Evolutionary Robotics [26], where a great range of robots/controllers must be constantly evaluated. One common approach is to perform the evaluation in simulation, and implement in real life the best performing solution. Likewise, we could sample the pdf of different robots in simulation, in order to estimate their pdfs in the real world. Of course, the accuracy of the pdf estimation would depend on the accuracy of the simulation environment.

In general, however, even without knowledge of the pdfs of the agents, this paper shows that a team composed by strong but very similar agents is not necessarily optimal. Hence, if a designer is not able to estimate the pdfs, she should at least evaluate the performance of diverse teams before picking only the strongest agents as the chosen team for a certain multi-agent application.

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

References

1. Marcolino, L.S., Jiang, A.X., Tambe, M.: Multi-agent team formation - diversity beats strength? In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence. (2013)
2. List, C., Goodin, R.E.: Epistemic democracy: Generalizing the condorcet jury theorem. *J. of Polit. Philosophy* **9** (2001) 277–306
3. Obata, T., Sugiyama, T., Hoki, K., Ito, T.: Consultation algorithm for computer shogi: Move decisions by majority. In: *Comp. and Games'10*. (2010) 156–165
4. Burke, R.: Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction* **12**(4) (November 2002) 331–370
5. Hong, L., Page, S.E., Baumol, W.J.: Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proc. of the Nat. Academy of Scienc. of the USA* **101**(46) (2004) 16385–16389
6. LiCalzi, M., Surucu, O.: The power of diversity over large solution spaces. *Manage. Sci.* **58**(7) (July 2012) 1408–1421
7. Condorcet, M.: *Essai sur l'application de l'analyse a la probabilité des décisions rendues a la pluralité des voix*. L'Imprimerie Royale (1785)
8. Young, H.P.: Optimal voting rules. *Journal of Economic Perspectives* **9**(1) (1995) 51–64
9. Conitzer, V., Sandholm, T.: Common voting rules as maximum likelihood estimators. In: *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (uai2005)*, Morgan Kaufmann Publishers (2005) 145–152

10. Xia, L.: Computational voting theory: game-theoretic and combinatorial aspects. PhD thesis, Duke University, Durham, NC, USA (2011)
11. He, L., Ioerger, T.R.: A quantitative model of capabilities in multi-agent systems. In Arabnia, H.R., Joshua, R., Mun, Y., eds.: Proceedings of the International Conference on Artificial Intelligence, IC-AI 03, June 23 - 26, 2003, Las Vegas, Nevada, USA, Volume 2, CSREA Press (2003) 730–736
12. Guttmann, C.: Making allocations collectively: Iterative group decision making under uncertainty. In Bergmann, R., Lindemann, G., Kirn, S., Pechoucek, M., eds.: Multiagent System Technologies, 6th German Conference, MATES 2008, Kaiserslautern, Germany, September 23-26, 2008. Proceedings. Volume 5244 of Lecture Notes in Computer Science., Springer (2008) 73–85
13. Liemhetcharat, S., Veloso, M.: Modeling and learning synergy for team formation with heterogeneous agents. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '12, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems (2012) 365–374
14. Agmon, N., Stone, P.: Leading ad hoc agents in joint action settings with multiple teammates. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1. AAMAS '12, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems (2012) 341–348
15. Gaston, M.E., desJardins, M.: Agent-organized networks for dynamic team formation. In: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems. AAMAS '05, New York, NY, USA, ACM (2005) 230–237
16. Matthews, T., Ramchurn, S.D., Chalkiadakis, G.: Competing with humans at fantasy football: Team formation in large partially-observable domains. In Hoffmann, J., Selman, B., eds.: AAAI, AAAI Press (2012)
17. Luan, S., Katsikopoulos, K., Reimer, T.: When does diversity trump ability (and vice versa) in group decision making? a simulation study. PLoS One **7**(2) (2012) e31043
18. Lakhani, K.R., Jeppesen, L.B., Lohse, P.A., Panetta, J.A.: The value of openness in scientific problem solving. HBS Working Paper (07-050) (2007)
19. Krause, S., James, R., Faria, J.J., Ruxton, G.D., Krause, J.: Swarm intelligence in humans: diversity can trump ability. Animal Behaviour **81**(5) (May 2011) 941–948
20. Braouezec, Y.: Committee, expert advice, and the weighted majority algorithm: An application to the pricing decision of a monopolist. Comput. Econ. **35**(3) (March 2010) 245–267
21. West, D., Dellana, S.: Diversity of ability and cognitive style for group decision processes. Inf. Sci. **179**(5) (February 2009) 542–558
22. Hellinger, E.: Neue begründung der theorie quadratischer formen von unendlichvielen veränderlichen. Journal für die reine und angewandte Mathematik **136** (1909) 210–271
23. McKelvey, R.D., McKelvey, R.D., McKelvey, R.D., Palfrey, T.R., Palfrey, T.R., Palfrey, T.R.: Quantal response equilibria for normal form games. In: Normal Form Games, Games and Economic Behavior. (1996) 6–38
24. Gelly, S., Wang, Y., Munos, R., Teytaud, O.: Modification of uct with patterns in monte-carlo go. Technical report (2006)
25. Pandanet: Introduction to Go. http://www.pandanet.co.jp/English/introduction_of_go/
26. Nolfi, S., Floreano, D.: Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines. MIT Press, Cambridge, MA (2001) 2001 (2nd print), 2000 (1st print).