Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

Collaboration in Social Problem-Solving: When Diversity Trumps Network Efficiency

Diego V. Noble, Marcelo O.R. Prates, Daniel S. Bossle and **Luís C. Lamb**
Institute of Informatics, Federal University of Rio Grande do Sul - Porto Alegre, 91501-970, Brazil
dvnoble@inf.ufrgs.br morprates@inf.ufrgs.br dsbossle@inf.ufrgs.br lamb@inf.ufrgs.br

Abstract

Recent studies have suggested that current agent-based models are not sufficiently sophisticated to reproduce results achieved by human collaborative learning and reasoning. Such studies suggest that humans are diverse and dynamic when solving problems socially. However, despite their relevance to problem-solving, these two behavioral features have not yet been fully investigated. In this paper we analyse a recent social problem-solving model and attempt to address its shortcomings. Specifically, we investigate the effects of separating exploitation from exploration in agent behaviors and explore the concept of diversity in such models. We found out that diverse populations outperform homogeneous ones in both efficient and inefficient networks. Finally, we show that agent diversity is more relevant than the strategic behavioral dynamics. This work contributes towards understanding the role of diverse and dynamic behaviors in social problem-solving as well as the advancement of state-of-art social problem-solving models.

1 Introduction

Collaborative problem-solving has been the subject of Artificial Intelligence research for over half a century (Newell, Shaw, and Simon 1959; Newell and Simon 1972; 1976; Simon 1990). Recent investigations in human computation, social computing and complex networks have also contributed to new results, tools and technologies for social and human problem-solving (Easley and Kleinberg 2010; Hogg and Huberman 2008; Kearns 2012; Law and von Ahn 2011). The analysis of collaboration has gained attention and several authors have indicated that it is an essential feature in problem-solving (Amir et al. 2013; Tabajara et al. 2013). By understanding the dimensions and patterns of collaborative human problem-solving, we can improve (human) organizations' performance, thus devising stronger ones (Araujo and Lamb 2008; March 1991; Page 2007).

In a recent study, (Mason and Watts 2012) have investigated relevant aspects of human collaborative learning in networks. They compared how humans and artificial agents solved the same problem in the same collaborative network context. Unsurprisingly, one of the findings was that humans behaved in a richer way than their artificial counterparts. Humans considerably outperformed artificial agents. They suggest that this result may indicate a lack of sophistication in current agent-based models, pointing to the fact that their agents — and the ones found in the literature — are "insufficiently sophisticated and heterogeneous to reflect real human responses to changing circumstances" (Mason and Watts 2012), lacking the dynamic behavior which humans presented in their problem-solving strategies. They go further and, accordingly to their findings: "the results of both artificial simulations and artificial experiments should be generalized with caution" (Mason and Watts 2012).

There is evidence that heterogeneity compensates for low ability in problem-solving (Marcolino et al. 2014). (Hong and Page 2004) state that *"diversity trumps ability"*. In their paper, they detail how a diverse pool of problem-solvers can outperform a homogeneous pool of high-ability agents. Although the authors did not address *dynamic* problem-solving agents (i.e. agents who change strategies throughout the solving process) the fact that human agents are likely to be dynamic *and* heterogeneous (Mason and Watts 2012) is an indicator that this question is still open.

This paper addresses the interplay between (1) collective diversity (heterogeneity), (2) dynamic strategic behavior, and (3) network efficiency, by extensive agent-based simulations. We show that the dissociation of exploitation from exploration can influence and change drastically what is an optimal individual strategy. More importantly, we found out that, from the three above parameters, collective diversity is shown more relevant to the collective performance than network efficiency and dynamic strategic behavior.

2 Preliminaries and Related Work

Computational social science is currently facing huge research interest (Lazer et al. 2009; Kearns 2012). Research in social problem-solving predates the recent ubiquity of social technologies. An early example is the work of (Clearwater, Huberman, and Hogg 1991; Clearwater, Hogg, and Huberman 1992), who discuss solving constraint-satisfaction problems by means of cooperative artificial agents.

More recently, (Hong and Page 2004) investigate how diverse agents teams can outperform high-ability ones. (Lazer and Friedman 2007) studied the dissemination of information in a network environment where artificial agents communicate with themselves in order to solve a particular complex problem [1]. They found that inefficient networks outperform efficient ones in the long-run. Network efficiency is measured as the average distance among the agents: an efficient network has small average distance between all node pairs, thus information flows faster.

(Mason and Watts 2012) experimented with both human and simulated agents to investigate collaborative learning in networks. They identified a gap between the results from human and computational experiments. They conjecture a lack of sophistication in current artificial agent-based models with respect to heterogeneity and dynamicity.

2.1 On Heterogeneity and Dynamicity

Our model builds upon the one employed by (Mason and Watts 2012), which in turn draws inspiration from (Lazer and Friedman 2007). We adapt Mason and Watts' model to support heterogeneity and dynamicity in the population, enabling different agents to have different search radii (heterogeneity) and also enabling them to change their radii during execution (dynamicity). The agents are tasked with finding the maximum of a particular real valued function. The precise meaning of the terms mentioned here and the details of the function solved by the agents are detailed below.

Search Radius: The search radius is an agent parameter responsible for establishing the boundaries of that agent's search space. Each agent's search space is bounded by a circle whose center is determined by the coordinates of that agent's current best solution and whose radius is that particular agent's search radius. It is only within the boundaries of this circle that the agent can search for new solutions in the *exploration* phase.

Myopic Search: A myopic search is a task performed by agents in the Lazer-Friedman (LF) model. It consists of selecting a point in the parameter space respecting the boundaries determined by the search radius.

Lazer-Friedman Model: (Lazer and Friedman 2007) is a model of social problem-solving in which agents are connected according to a particular network topology. The problem is solved after a number of iterations. At each iteration, every agent updates its own state; if the agent in question (focal agent) holds the best solution it performs *myopic search*; otherwise it copies the solution from the best neighbor agent.

Heterogeneity (diversity): An agent population is *heterogeneous* when agents have different search radii. A population is *homogeneous* when all agents have the same search radii.

Dynamicity: When agents are able to change their search radius when solving a problem, they are dynamic. In a static population all agents are endowed with predetermined search radii throughout the entire problem-solving task.

2.2 On Annular Myopic Searches

In a myopic search, agents select a point in the parameter space respecting the boundaries determined by their search radii. As our results will show, this approach can be problematic: the fact that some agents have bigger circles than others puts the second group in an unfair disadvantage towards the first one, as its agents have their search space reduced in comparison. The problem with this disadvantage is that it generates a bias toward homogeneous populations of high-ability agents and against populations of diverse ones. It is always possible for a homogeneous population to outperform a heterogeneous one - the only necessary condition is that the global radius of the homogeneous population is greater than every single radii of the heterogeneous one.

This conclusion is logical, but is nevertheless inconsistent with the generally accepted idea that, at least to some extent, *diversity trumps ability*. It is particularly inconsistent with the findings of (Hong and Page 2004) who show that high ability agents are often outperformed by diverse ones. Naturally, the inconsistency is not due to the fact that diversity actually does not trump ability - it does - but rather to the fact that using search radii to determine circular search spaces creates a bias against heterogeneous populations. This happens because circular search spaces fail to include the exploration-exploitation dichotomy in the problem-solving task, as agents with larger circles are just as capable of exploiting their current solutions as agents with smaller circles, but additionally capable of exploring new solutions in a larger area. In short, although we can create heterogeneous populations by diversifying their radii, we cannot create a heterogeneous population in which agents have varied inclinations towards exploration and exploitation - we can only diversify their search.

A strategy to solve this problem is to restrict the search space of exploration-inclined agents to outside their close neighborhood, making it impossible for them to exploit their current solutions. Exploitation-inclined agents should also have their search space restricted, but this time to include *only* their close neighborhood. This strategy can be implemented by replacing the circles in the myopic search by rings. Thus, when we enlarge a particular agent's search radius, we are not only enlarging the ring's outer circle but the ring's inner circle as well. This forces this agent to explore in a region that becomes increasingly distant of the agent's current solution as its search radius grows. By manipulating an agent's search radius, we can place this agent wherever we want in the exploration-exploitation continuum. We refer to our model as the ***Ring Model***.

3 Experimental Setup

In our experiments, we tested the Mason & Watts instance of the LF model (Mason and Watts 2012) as well as our proposed *Ring Model* and its variations. Although these variations implement different logical rules to solve a problem, their algorithmic structure is the same - as described in Algorithm 1.

Each agent is instantiated with a valid solution (line 1) then it is embedded in a social network (line 2). Each social tie establishes that an agent will perform bidirectional communication with another agent to whom it is connected. After these two stages, each agent performs the search. The agent samples a new solution according to the procedure

[1] They used the maximization of a rugged continuous function generated by NK-model (Kauffman and Weinberger 1989).

Algorithm 1: General algorithm structure used by all models.

Input : $Population_{size}$, $stop\_condition$, $Problem_{size}$
Output: $Solution_{best}$

```
   // Initialize the population with random solutions.
1  Population ← InitPopulation (Population_size, Problem_size)
   // Establish social ties among individuals.
2  GenerateNetwork (Population)
3  repeat
4      foreach node ∈ Population do
           // Execute the model search procedures.
5          NextState (node, Problem_size)
6      foreach node ∈ Population do
           // Decide to accept or not the new state.
7          UpdateState (node, Problem_size)
   until stop_condition
8  Solution_best ← RetrieveBestSolution (Population)
9  return Solution_best
```

Algorithm 2: Standard Ring Model procedure.

```
1  if node_solution < max(Neigbors' solutions) then
       // Local search procedure (or myopic search).
2      repeat
3          r ← Random (low = max(node_radius − 1, 0) high = radius)
4          angle ← Random (low = 0, high = 2 ∗ π)
5          x ← RandomInt (r∗cos(angle)+node_x)
6          y ← RandomInt (r∗sin(angle)+node_y)
       until Until a valid solution is found
   else
       // Copy procedure.
7      (x,y) ← RetrieveBestSolution (node_Neighbors)
8  node_nextSolution ← (x,y)
```

NextState (line 5). This procedure is where each model implements its own search rules (i.e. if the sampling will be done within a full circle or within a ring shaped area). If the new solution improves the current one, then the agent accepts it, otherwise it will hold the old solution (line 7). An iteration is complete after all agents have their search turn. The process is repeated until there is no more iterations to be done. We limited the number of iterations to one hundred.

Algorithm 2 details how each individual executes the *NextState* procedure in our Ring Model. The first step is to detect if the solver (agent) holds the best solution among its neighbors (line 1). If this is the case, then the agent performs the myopic search procedure, that is randomly selecting the coordinates of the next solution to be sampled. The coordinates in lines 5 and 6 are bound by a ring shaped area with a constant radius and constant thickness (lines 2 to 6). If the agent is not among the best in its neighborhood, the agent copies both coordinates from her neighbor (line 7). A solution is valid if it is within the search space coordinates.

The Ring Model can be divided into four types:

Static Homogeneous — All agents have the same fixed radius throughout the search process.
Static Heterogeneous — Though fixed, there may be diverse radii within the same population.
Dynamic Social — Agents copy not only their neighbor's solutions, but their radius at each iteration.
Dynamic Random — At each iteration, the agent chooses a random radius from a limited pool of choices.
The problem instances were generated using a similar process to that of (Mason and Watts 2012). The process generates a continuous bi-dimensional problem space which contains noise - generated using the Perlin noise method (Perlin 1985) - and a single signal - generated by summing a Gaussian to this noise.

Figure 1 depicts one such instance. The size of the search spaces is fixed at 200×200 integer points - i.e. each solution is a pair (x, y) where x and y are integers within the range $[0, 200)$. Each combination of (x, y) represents a solution which has a respective score and the problem consists basically of finding the pair that maximizes the function output. The coordinates of the signal are chosen randomly and the parameters to generate problem instances are: $\omega = [2, 3, 4, 5, 6]$, $\sigma = 9$, and persistence parameter $\varpi = 0.7$. The best solution (signal) is the black region while the noise is depicted as the light grey region. We generated a new instance for each independent trial and used 1000 independent trials. Moreover, our search space is larger and more rugged than the search space employed by Mason & Watts. This leads to a problem that is harder to solve. The peak (best solution) has score of a 100.0. Successful trials are the ones in which any agent achieve a score equal to or higher than 95.0. Moreover, we discarded trials where any agent's initial score score was higher than or equal to 55.0 in order to avoid the case where an agent is positioned too close to the peak (best solution).

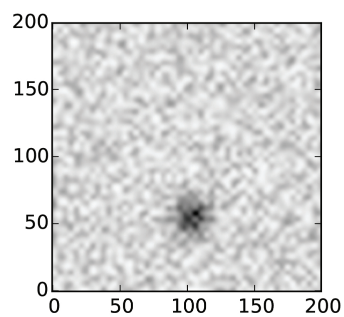

Figure 1: Sample problem instance: The peak is represented in black; the noise, in light grey.

4 Dissociating Exploitation from Exploration

One of our aims is to compare the effects of exploration and exploitation in agent behaviors in social problem-solving systems. We also analyse when such effects are mutually exclusive. This means that we consider the effects of using current information to solve a problem, while keeping or not the chance of finding rather novel solutions.

In this experiment, we compare the LF model instantiated by (Mason and Watts 2012), with our proposed Ring Model. The probability of finding the peak for both models is shown in Figures 2a and 2b. The main parameters used are the radius and the population size. We considered the first 32 rounds for this experiment and arranged individuals in a bi-dimensional periodic grid network where all agents have the same degree and are at the same average distance from any other node. By controlling the network, we remove any effect that a network position may have on the agent behavior (Kearns, Suri, and Montfort 2006; Judd, Kearns, and Vorobeychik 2010). Vertical bars represent 99% confidence intervals using the Wilson score interval method. Both models presented different behaviors

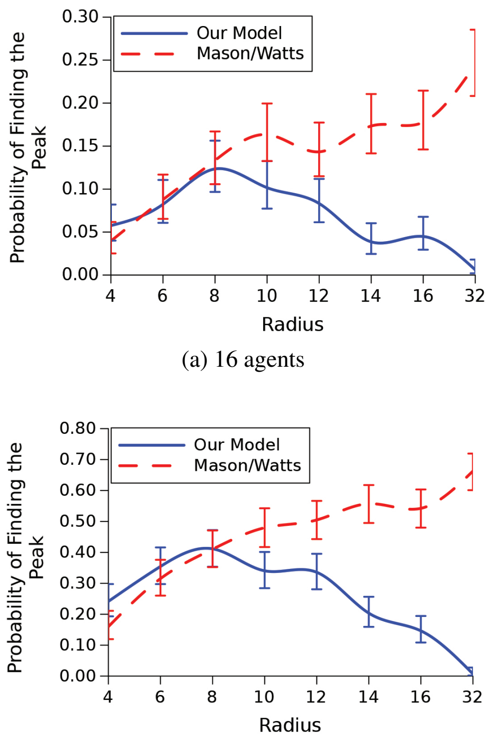

(a) 16 agents

(b) 64 agents

Figure 2: Probability of finding the peak for various radius sizes in our Ring Model and in the Mason & Watts model.

when individuals have bigger radii. The optimal behavior, w.r.t. radius size, is around 8 (4% of the grid's size) in the Ring Model, but in the Mason & Watts model, the bigger the radius the better. Thus, even if these models do not significantly differ for radius equal to or less than 8, for bigger radii they clearly depart. This trend is valid for both population sizes. Therefore, when one has a homogeneous agents' population that have traits associated with exploratory behavior, one must pay attention to which kind of exploratory behavior the agents employ. If exploration and exploitation are packed together (as in the Mason & Watts' models), the highest the exploratory bias, the better. However, if this is not the case, as it is in the Ring Model, that exploratory and exploitative behaviors co-exist separately; one must fine tune the agent group to the optimal strategy. We were able to reproduce the effect we described before in the Mason & Watts's version of the LF model: as the individual radius increased, so did the collective performance.

5 Diversity, Strategy and Network Efficiency

Until now, we experimented with the Ring Model and a homogeneous population where every agent held the same radius during the entire solving process. From now on, we will incorporate diversity into the Ring Model, with both static and dynamic diversity. Static diversity is characterized by agents starting with different radii and holding such radii until the end of the simulation, while in dynamic diversity they start with different radii, but can change their own radius during the simulation. Recall that the Ring Model with static diversity is known as the Static Heterogeneous Model. In this model, agents start with a radius of 4, 8, 16, 24, or 32 units. Each radius has the same chance of being chosen (the distribution is uniform).

Dynamic diversity invites an interesting question: *how can strategies be changed on an individual basis?* To answer this question we designed two versions of the Ring Model with dynamic diversity. The first is called the *Social Model*, where agents copy not only the solution of their best neighbor, but also their radius at each iteration; the second one is the *Random Model*, where agents always change their radius randomly after performing search procedures. In both models, agents start with a random radius. In the random model, the radii to be chosen are from the same radii pool used in the Static Heterogeneous model.

Figure 3 depicts the average score of the population over all trials. We consider 4 models: (1) the homogeneous Ring Model with radius 8, (2) the Static Heterogeneous Model, (3) the Random Model, and (4) the Social Model. The average score includes all trials, whether or not the trial is successful. We isolated the effect of the individual position in the network by using a bi-dimensional periodic grid topology. For both population sizes, we observed that the Random Model and the Heterogeneous Model presented the same average behavior. The difference between the Heterogeneous Model and the Random Model is not significant for population sizes of 16 and 64 (p=0.2 and p=0.9, respectively). Thus, random dynamic diversity does not statistically differ from static diversity in this context. We can also see that the Social Model never outperformed the Static Heterogeneous, nor the Random Model; it is not significantly better than the Static Homogeneous model for a bigger population (p=0.57). Even if agents are dynamic, and diverse, the Social Model is not better than the Static Homogeneous Model. When we remove the influence of network topology our results suggest that diversity could be more relevant than individual strategy dynamics which supports results of (Judd, Kearns, and Vorobeychik 2010).

We then test how network efficiency impacts each model. We define *network efficiency the as the speed that solutions flow through the network*, which is related to the average shortest path between any two network nodes (Mason and Watts 2012). An efficient network has small average path

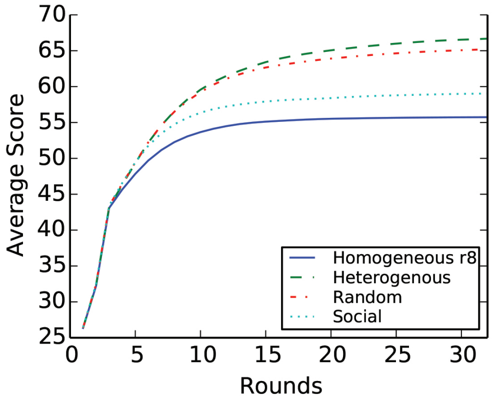

(a) 16 agents

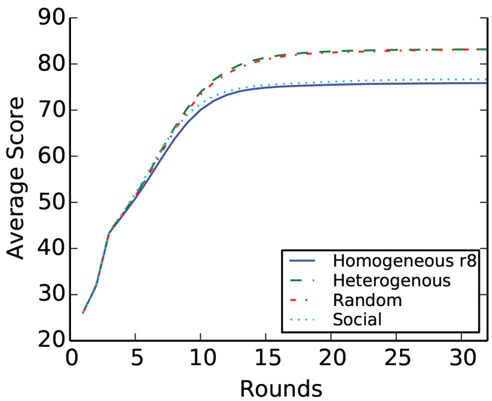

(b) 64 agents

Figure 3: Average scores over time for different population sizes and models.

lengths while an inefficient has large ones; i.e. an efficient network can be seen as a decentralized network where a good solution spreads faster on average when compared to a centralized network where the average distance between peripheral and central agents tend to be higher. The inefficient (centralized) network was built using the generative model of (Barabási and Albert 1999). This model produces a network topology with a power-law distribution. This distribution guarantees a centralization of the network as a minority of agents will hold most connections. The periodic grid topology was used as the efficient (decentralized) network.

Table 1 presents data from the experiments for all models in efficient and inefficient networks. Values between parenthesis represent 99% confidence interval using the Wilson score. We used the population size 64. We observe that a small population of 16 agents would present a less clear difference between an efficient and an inefficient network. We measured these probabilities after the 32nd round, therefore if any agent found the best solution (peak), we considered it a successful trial.

The difference among the above models was much stronger and clear than the difference regarding network efficiency. With the exception of the Social Model, all models performed slightly better in efficient (decentralized) networks when compared to inefficient (centralized) networks. However, the likelihood of finding the best solution varied considerably among the models. The Dynamic Random and the Static Heterogeneous models presented the highest odds of finding the solution. The small difference between these models is evidence that static diversity is as effective as dynamic diversity. Among the Static Homogeneous versions, the Ring Model with radius 8 was the best ranked model in both networks. In this case, the difference between network efficiency accounted for approximately 20%, while between the Ring Model with radius 16, the difference was approximately 64%. The model with radius 32 presented the worst results for both networks. The Social Model was the only one that took advantage of an inefficient network; its performance increased by approximately 17% when compared to the efficient network. These results provide evidence - by means of artificial agents simulations - that individual behavior may have a stronger influence in the collective outcome than network topologies themselves. This evidence, as previous researchers have shown, reinforce that "the intrinsic behavioral traits of human actors may play a stronger role in their relative influence than purely structural properties of their network position" (Judd, Kearns, and Vorobeychik 2010).

Model	Network	
	Efficient	Inefficient
Homogeneous radius 8	0.41 ($\pm$ 0.06)	0.33 ($\pm$ 0.06)
Homogeneous radius 16	0.15 ($\pm$ 0.04)	0.13 ($\pm$ 0.04)
Homogeneous radius 32	0.01 ($\pm$ 0.01)	0.03 ($\pm$ 0.02)
Static Heterogeneous	0.62 ($\pm$ 0.06)	0.59 ($\pm$ 0.06)
Dynamic Random	0.62 ($\pm$ 0.06)	0.54 ($\pm$ 0.06)
Dynamic Social	0.23 ($\pm$ 0.05)	0.27 ($\pm$ 0.05)

Table 1: Probabilities of finding the peak for efficient (decentralized) and inefficient (centralized) networks.

In the sequel, we tested if the most central agent would reach the best solution earlier than the median agent or the least central one. In order to do so, we retrieved the round that each of the above agents found the solution, from all independent trials. One could expect central agents to find the best solution before the peripheral ones. Results are depicted in Figure 4. We used betweenness (Borgatti and Everett 2006) as a centrality metric and used the Barabasi-Albert (Barabási and Albert 1999) generative model to generate the networks. Results are depicted in Figure 4. For this problem, an agent's individual position in a network was irrelevant. On average, the least central agent, the median and the most central agent achieved the best solutions at the same round as seen in Figure 4.

6 Analysis and Discussion of Results

In homogeneous populations, all agents are restricted to a fixed search radius. As a result, all of them see the search space through the lens of exactly the same ring. For exam-

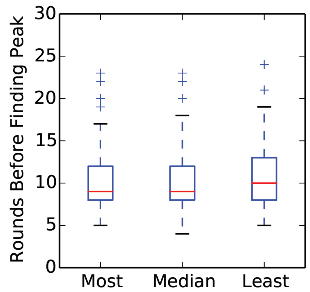

(a) Ring 8 model

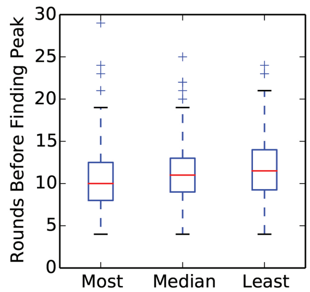

(b) Random model

Figure 4: Rounds before finding the peak (solution) for the most, median, and least central node in: (a) Static Homogeneous with radius 8, and the (b) Dynamic Random model.

ple, a possible setup for an homogeneous population consists of all agents having a search ring delimited by an outer circle of radius 5 and an inner circle of radius 1. We could also envision a setup in which all agents have an outer radius of 10 and an inner radius of 8.7. Although the rings used by agents in both examples would have approximately the same area (24π and 24.3π), they would render two completely different strategies. The first setup would render the population highly exploitation-oriented, as each agent would be able to see its surroundings in great detail, while not being granted access to further areas which would usually be searched by exploration-oriented agents. In contrast, the second setup above (radius 10 and inner radius of 8.7) would render the population highly exploration-oriented, being unable to exploit their current solution and thus would be forced to search for other, far away solutions.

The search space we used in our experiments was a bi-dimensional Gaussian function with additional random noise. Let us consider an one-dimensional analogue of this search space. For simplicity, let us also assume we have a sinusoidal wave of fixed frequency instead of random noise. This search space has a global maximum near $x = 0$, but several well-distributed local maxima. As a result, agents that rely solely on exploitation will have a hard time finding the global maximum, often getting stuck at local ones. As each agent's ring is limited (agents cannot see through the ring's cavity), agents will sometimes be unable to grasp the general behavior of the curve in a given point, and would refuse to accept a solution that would bring them closer to the global maximum, even though they can see it with their annular vision.

In a homogeneous population, the situation described above is bound to happen very often. As all agents have the same search radius, the population as a whole will not be able to overcome the pitfalls of their particular search ring, with its agents finding themselves stuck on local maxima a lot more often than they should. Conversely, in a heterogeneous population, the diversity of shapes agents' rings can take assures that no pitfall is more common than any other. As a result, the probability that at least one search ring will meet the necessary criteria to get its agent climbing the search space without getting stuck in any point is maximized. Therefore the social dynamics of the agent model can then spread this agent's solution throughout the entire population.

7 Conclusions

Collaboration is a relevant research issue in both human and AI problem-solving. Mason and Watts' results suggest that artificial agents were vastly outperformed by human subjects (Mason and Watts 2012). Our initial hypothesis was that the lack of diversity in artificial agent populations was detrimental to their performance. As previously suggested in the literature (Hong and Page 2004; Marcolino et al. 2014), diversity (or heterogeneity) in artificial agent models was thought to reduce the gap between human and artificial agent collaborative behaviors. In our investigation, we conducted a series of experiments testing the performance of homogeneous and heterogeneous populations. Our results not only clearly pointed out that heterogeneous populations have a significant advantage when compared to homogeneous counterparts, but also shed new light on a number of interesting proprieties regarding artificial agent populations and their diversity. In particular, we found out that the performance of homogeneous populations decrease with large radii.

We have drawn qualitative explanations of the rules we discovered to dictate the performance of homogeneous and heterogeneous agent populations. We have also analysed the relationship between our findings and the work of others. We have shown how our results about the benefits of heterogeneity are in conformity to those of (Hong and Page 2004). Ultimately, we added supportive evidence that heterogeneous (diverse) populations are superior, performance-wise, to homogeneous ones and that diversity is more relevant to collaborative problem-solving than dynamic behavior and network efficiency.

Acknowledgments

This work is partly supported by the Brazilian Agencies CAPES and CNPq.

References

Amir, O.; Shahar, Y.; Gal, Y.; and Ilani, L. 2013. On the verification complexity of group decision-making tasks. In *First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013*. AAAI Press.

Araujo, R. M., and Lamb, L. C. 2008. Memetic networks: Analyzing the effects of network properties in multi-agent performance. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-08)*, 3–8. AAAI Press.

Barabási, A., and Albert, R. 1999. Emergence of Scaling in Random Networks. *Science* 286(5439):509–512.

Borgatti, S. P., and Everett, M. G. 2006. A Graph-theoretic perspective on centrality. *Social Networks* 28(4):466–484.

Clearwater, S. H.; Hogg, T.; and Huberman, B. A. 1992. Cooperative problem solving. In *Computation: The Micro and the Macro View*, 33–70. World Scientific.

Clearwater, S. H.; Huberman, B. A.; and Hogg, T. 1991. Cooperative solution of constraint satisfaction problems. *Science* 254(5035):1181–3.

Easley, D. A., and Kleinberg, J. M. 2010. *Networks, Crowds, and Markets - Reasoning About a Highly Connected World*. Cambridge University Press.

Hogg, T., and Huberman, B. A. 2008. Solving the organizational free riding problem with social networks. In *AAAI Spring Symposium: Social Information Processing*, 24–29. AAAI Press.

Hong, L., and Page, S. E. 2004. 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 (PNAS)* 101(46):16385–16389.

Judd, S.; Kearns, M.; and Vorobeychik, Y. 2010. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 107(34):14978–14982.

Kauffman, S. A., and Weinberger, E. D. 1989. The nk model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology* 141(2):211 – 245.

Kearns, M.; Suri, S.; and Montfort, N. 2006. An experimental study of the coloring problem on human subject networks. *Science* 313(5788):824–7.

Kearns, M. 2012. Experiments in social computation. *Commun. ACM* 55(10):56–67.

Law, E., and von Ahn, L. 2011. *Human Computation*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

Lazer, D., and Friedman, A. 2007. The Network Structure of Exploration and Exploitation. *Administrative Science Quarterly* 52(4):667–694.

Lazer, D.; Pentland, A.; Adamic, L.; Aral, S.; Barabasi, A.-L.; Brewer, D.; Christakis, N.; Contractor, N.; Fowler, J.; Gutmann, M.; Jebara, T.; King, G.; Macy, M.; Roy, D.; and Alstyne, M. V. 2009. Social science: Computational social science. *Science* 323(5915):721–723.

March, J. G. 1991. Exploration and Exploitation in Organizational Learning. *Organization Science* 2(1):71–87.

Marcolino, L. S.; Xu, H.; Jiang, A. X.; Tambe, M.; and Bowring, E. 2014. Give a hard problem to a diverse team: Exploring large action spaces. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI-14)*. AAAI Press.

Mason, W., and Watts, D. J. 2012. Collaborative learning in networks. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 109(3):764–7699.

Newell, A., and Simon, H. 1972. *Human problem solving*. Prentice-Hall.

Newell, A., and Simon, H. A. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the ACM* 19(3):113–126.

Newell, A.; Shaw, J. C.; and Simon, H. A. 1959. Report on a general problem-solving program. In *IFIP Congress*, 256–264.

Page, S. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press.

Perlin, K. 1985. An image synthesizer. *SIGGRAPH Comput. Graph.* 19(3):287–296.

Simon, H. A. 1990. Invariants of Human Behavior. *Annual Review of Psychology* 41(1):1–20.

Tabajara, L. M.; Prates, M. O. R.; Noble, D. V.; and Lamb, L. C. 2013. Leveraging collaboration: A methodology for the design of social problem-solving systems. In Hartman, B., and Horvitz, E., eds., *First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013*. AAAI Press.