The Great Melting Pot: generating diversity by combining solutions across a global population

EDUARDO J. IZQUIERDO, CHELSEA M. CAMPBELL and ROBERT L. GOLDSTONE, Indiana University Bloomington

1. INTRODUCTION

One of the major ways that people engage in adaptive problem solving is by copying or imitating the solutions of others. Imitation saves an individual time and mitigates potential risks from individual trial-and-error learning. When an individual finds a neighbor with a better solution than theirs, copying their entire solution guarantees an improvement over the individual's current condition. However, this reduces the diversity of solutions in the group and can lead the group to getting stuck in a local optima. One alternative is to copy the neighbor's solution only partially, although this comes at a risk for the individual. Mixing two solutions may or may not lead to an improvement over their previous solution, but mixing has the potential to allow the group to explore entirely new areas of solution space. So, although partial copying comes at a cost to the individual, under what conditions does it benefit the group? In the current research, we are interested in the consequences for the group when its members engage in social learning strategies with different tendencies to copy entire or partial solutions, with different network topologies that affect the neighbors' solutions visible to each member, and with different complexities of search tasks.

2. RELATED WORK

A number of studies have used computer simulations to address how different social learning strategies perform in tasks of varying levels of difficulty. These studies have focused primarily on the effect that the structure of the network of interactions has on the average performance of the group (Lazer & Friedman, 2007; Fang et al., 2010; Mason & Watts, 2012; Derex et al., 2015; Barkoczi & Galesic, 2016). Most of these studies have found that when the sharing strategy is to copy the best solution in the individual's neighborhood (henceforth 'best member' strategy), locally-connected groups outperform globally-connected ones (Lazer & Friedman, 2007; Fang et al., 2010; Derex et al., 2015; Barkoczi & Galesic, 2016).; others have found that when the sharing strategy is to copy the most common solution in the individual's neighborhood (henceforth 'conformity' strategy), the pattern is reversed: globally-connected groups outperform locally-connected ones (Mason & Watts, 2012; Barkoczi & Galesic, 2016).

3. MODEL

We develop a simulation model of problem solving, where a group of individuals are repeatedly searching for solutions.

3.1 Problem spaces

We follow previous work in modelling this problem as a search on NK tunably rugged landscapes (Kauffman, 1995). In Kauffman's formulation, N is the number of traits and K specifies the extent to which fitness is determined by interactions between traits. For our purposes, N is associated with the size of the problem, and K is associated with the degree of difficulty of the problem. A problem space with no trait interactions (K=0) has a single global optimum. As K increases, the number of local optima increases and the problem space becomes more rugged. The main challenge for a group is to find the better solutions without getting stuck in local optima.

3.2 Social learning strategies

At any given point in time, each individual in the group has a solution in the NK space. Each solution has a score associated with it. The group starts with random initial solutions and each individual in the group has 200 trials to increase their score. During each trial, every individual gets to share information with one of its neighbors. A neighbor is picked randomly, and if the neighbor has a better scoring solution than the individual, the individual adopts the neighbor's solution. If that neighbor's solution is worse, then the individual learns on its own through trial-and-error. An individual does this by altering one of the bits of its own solution. If the change leads to a better score, then the individual adopts that solution. If the change produces a worse-scoring solution, then the individual discards it and reverts back to its original solution. Our model has one primary difference to the social learning strategies from previous modeling work. Most previous studies have considered that when a neighbor has a better strategy the individual adopts 100% of the others' strategy. We also consider cases where the individual adopts the better individuals' solution only partially (i.e., 50% of the solution).

3.3 Network of collaborations

Individuals in a group are connected to each other through a network of collaborations. Different network topologies affect the solutions that are visible to each member. In this work, we report on two extremes: globally- and locally-connected groups. In globally-connected groups, every member of the group can copy solutions from every other member of the group. In locally-connected groups, the individuals are geographically distributed on a 1D ring and each individual only has access to solutions from their immediate neighbors.

3.4 Group performance

Most studies have only considered the average group performance. We examine the performance of best member, the average for the group, and the number of groups to reach the global optima. Because NK problem spaces are notoriously diverse, results reported here are compared using the same set of 1000 NK problem spaces and the same set of initial starting solutions of the group. Also, for each NK problem space, the score is normalized to the min and max possible score for that space.

4. RESULTS

We ran simulations in problem spaces of size N=15, with K varying between 0 and 14, for globally- and locally-connected groups of 100 individuals, and for social learning strategies that involved copying entire solutions or only partially copying solutions. Each group was given 200 learning/copying trials to find solutions. Each condition was ran on 1000 different problem spaces and starting solutions. Our simulations revealed three key insights. First, partial sharing benefits from globally-connected groups on complex tasks. Consistent with previous studies, we find that when sharing involves copying entire solutions, locally-connected groups outperform globally-connected groups in complex search spaces. However, when copying is partial, the pattern is reversed: globally-connected groups outperform locally-connected ones. Second, when success is measured by average group performance, groups with partial sharing outperform both locally- and globally-connected groups with full sharing. Third, globally-connected groups whose members share partial solutions outperform even groups of pure individual learners at finding the optimum solution to a problem on complex tasks. This is in contrast to strategies that involve copying entire solutions, which are always outperformed by groups of pure individual learners.

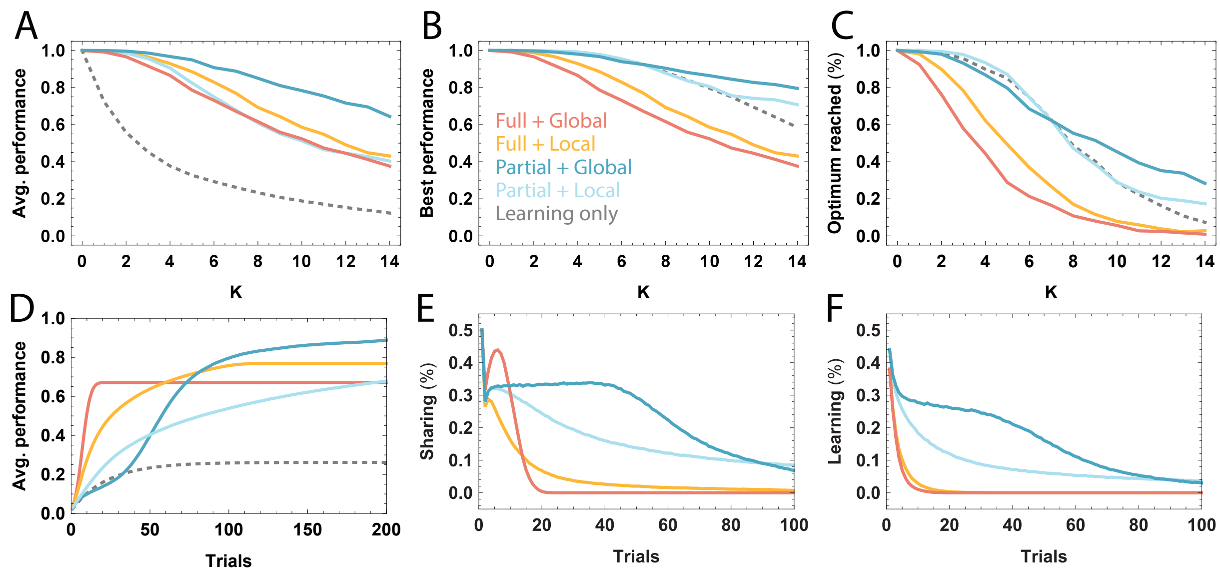

Fig. 1. Simulation results with N=15, K=[0,14], group size 100, 200 learning/copying trials, and 1000 repetitions. Performance of five conditions as a function of ruggedness [A-C]. Full-copying in globally- and locally-connected networks in red and orange, respectively. Partial-copying in globally- and locally connected networks in dark and light cyan, respectively. Gray dashed trace depicts a pure individual learning strategy. Group learning dynamics over time for K=7 problem spaces [D-F]. Average performance of the groups over time [D]. Proportion of individuals in the group that shared successfully with a neighbor over time [E]. Proportion of individuals in the group that successfully learned individually over time [F].

5. DISCUSSION

Diversity is good for broad exploration. Convergence is good for propagating good solutions. The ideal amount of diversity depends on the problem complexity (e.g. K in the NK framework). As complexity increases, populations with greater diversity perform best. For example, the 'best member' strategy with 'full' copying benefits from locally- rather than globally-connected members because it slows down the communication of best solutions found thus far, preserving diversity in the group. The 'conformity' strategy prolongs diversity by focusing on individual learning initially (while all the solutions are mostly different) and then as solutions converge, gradually shifting to sharing the best common solutions. However, this social learning strategy deteriorates quickly as the size of the problem space increases, effectively turning into a pure individual learning strategy. We show that a partial-sharing strategy can also provide diversity. The diversity introduced by partial sharing, however, is generated through the combination of different solutions in a problem space. Thus, diversity is not merely preserved but actively *generated* under this strategy. This explains why the partial-sharing model does not benefit from locally-connected networks. Groups that are sharing partially benefit from combining entirely different solutions. Our work suggests that what is good for an individual is not necessarily good for the group. When copying involves adopting the exact solution of a better-performing neighbor, the individual is guaranteed an improvement in score. This is not the case when copying involves partial adoption: mixing two solutions may or may not lead to an improvement in score. So although partial sharing comes at a cost to the individual, we show that there are clear benefits for the group in terms of performance, and that the benefits increase proportional to the difficulty of the task.

REFERENCES

Lazer, D., & Friedman, A. (2007). The Network Structure of Exploration and Exploitation. Administrative Science Quarterly, 52(4), 667694. 10.2189/asqu.52.4.667

Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. Nature Communications, 7. 10.1038/ncomms13109

Fang, C., Lee, J., & Schilling, M. (2010). Balancing Exploration and Exploitation Through Structural Design: The Isolation of Subgroups and Organizational Learning. Organization Science. 21. 625-642. 10.1287/orsc.1090.0468.

Derex, M., Feron, R., Godelle, B., & Raymond, M. (2015) Social learning and the replication process: an experimental investigation. Proceedings of the Royal Society B: Biological Sciences, 282(1808). 10.1098/rspb.2015.0719

Mason, W., Watts, D.J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*. 109(3)764-769. 10.1073/pnas.1110069108

Kauffman, S. (1995) At Home in the Universe: The Search for Laws of Self-Organization and Complexity. New York: Oxford University Press.