

Optimizing the efficiency of collective decision making in groups

Malgorzata Turalska*, Rosie Lickorish†, Geeth De Mel†, Liam Turner‡, Roger M. Whitaker‡

*Network Science Division, CCDC Army Research Laboratory, Adelphi, MD, USA

†IBM Research Europe, Hursley, UK

‡Crime and Security Research Institute, Cardiff University, Cardiff, UK

Abstract—The complexity of modern military operations create a demand for efficient collaborative decision making and problem solving. Additionally, as military units operate in increasingly dynamic environments, the ability to respond to changing circumstances becomes paramount for mission success. An effective response rests on correct dissemination and transfer of information across the command and control structure, and thus is critically linked to the network of human interactions.

In this paper, we take an agent-based modeling approach to collective problem solving. We investigate three key factors affecting the performance in collaborative environments: (1) the structure of network used to share information between agents, (2) the search strategies adopted by agents, and (3) the complexity of problems facing the group. In particular we study how the trade-off between exploitation of known solutions and exploration for novel ones is related to the efficiency of collective search. Additionally we consider the role of agent behavior: propensity for risk-taking and trustworthiness, as well as the dynamic nature of social connections. Finally, we outline the directions for future work regarding the efficiency of problem solving on military-like command and control structures.

Index Terms—collective decision making, NK-landscapes, information spread, collective intelligence, problem solving

I. INTRODUCTION

Problems encountered in business, engineering and governance are increasingly complex in the sense that they require optimization of multiple variables, where the dimensionality and interdependence between problem composites extends beyond cognitive abilities of individual decision makers [11], [40]. Therefore, with the increasing scale of tasks facing managers, engineers and planners, collaboration and teamwork become primary mechanisms of decision making, problem solving and generation of knowledge and innovations [41]. Teams which engage in collaborative problem solving by harnessing diverse and complementary skills [13], [15] are critical in addressing complex and dynamic tasks that require efficient processing of distributed and networked information [23].

Theoretical and experimental work on collective intelligence distinguishes two main paths to the acquisition of knowledge [4], [12], [25]. Individuals can innovate by searching for new solutions through trial and error, testing various courses of action in isolation from others, or alternatively, they can copy existing solutions by imitating other individuals [24]. Trial and error approach can lead to potentially novel solutions, however due to bounded rationality of decision makers, this exploration

is myopic in scope and restricted to testing only few dimensions of the current solution [19], [20]. Imitation, on the other hand, is a simple cognitive strategy that can lead to exceptionally good outcomes [16], [25], [28], [31], however it cannot produce new information by itself. Additionally, social learning can be error prone [22], leading groups to acquire inappropriate or outdated information in nonuniform and changing environments [7], [30]. Altogether, cognitive learning in groups can be seen as the outcome of both individual learning and the spread of social information through imitation [10], [14]. Different balance of these two strategies can lead to very different levels of group performance [3], [13], [33].

In many regards modern military operations experience analogous problem solving challenges as the ones encountered by companies and organizations. As in civilian institutions, the ability to efficiently collect and disseminate information and share it through the command and control (C2) structure is paramount for success of military enterprises. However, the unique objectives facing the military result in a distinct set of issues related to collaborative problem solving. The extreme uncertainty of military mission space, stemming from high intensity of combat operations, coupled with a wide spectrum of scenarios, ranging from humanitarian relief operations to complex coalition efforts, creates challenges not experienced

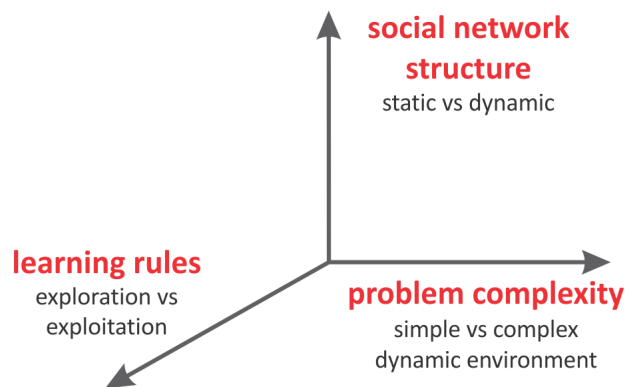


Figure 1. Accuracy and success of collaborative problem solving are functions of three fundamental factors: the structure of the social network through which information is exchanged, the complexity of problem that requires solution and strategies adopted in searching for the solution.

by modern companies. The dynamic and nonlinear character of battle space demands increased C2 agility obtained through augmented information sharing and raised shared awareness [1].

The key question is what types of command structures are most suitable for achieving the required C2 capabilities. In this paper, we adopt an agent-based model to investigate three key aspects of collective problem solving (Fig. 1) in an attempt at understanding this issue. We study how: (1) the structure of networks used to share information between agents, (2) the search strategies adopted by agents, and (3) the complexity of problems facing the group affect the performance of teams working in collaborative environments. We demonstrate that different levels of performance can be explained by varying exploration-exploitation patterns, and those in return translate into amount of information gathered during the collaborative problem solving. The information-based approach provides a framework for quantifying the impact of various configurations of collective problem solving environments and allows us to formulate directions of further work.

II. METHODS

A. Fitness landscape

A common way to represent complexity of problems is with multidimensional search spaces, where any possible combination of choices is represented by a point in the landscape, a height of which denotes the performance of that combination [20], [35], [36]. The process of problem solving is then seen as sequential search, where individuals are characterized by limited rationality, and thus are not able to identify optimal solutions in a single try [37].

The NK model [17], [18] is a fundamental computational implementation of fitness landscapes that has been extensively used to study complex problem representation in management research [20], [35], [36]. The model represents a complex problem as one consisting of N variables where each component can take two possible values. Thus, the fitness landscape consists of 2^N possible policy choices, with the overall problem space characterized by a vector $\mathbf{x} = (x_1, x_2, \dots, x_N)$, where each x_i

takes on the value of 0 or 1. The parameter $K = 0, \dots, N - 1$, determines the level of interdependence between components and influences strongly the number of local maxima in the landscape.

To each string \mathbf{x} representing a unique solution we associate a fitness value $\mathcal{F}(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x})$, where f_i is the contribution of component i to the fitness of string \mathbf{x} . The local fitness f_i depends on the state x_i as well as in the states of K distinct randomly chosen components. The values of each f_i are assigned randomly from a uniform distribution in the unit interval and because of the randomness of f_i , we can guarantee that $\mathcal{F}(\mathbf{x}) \in (0, 1)$ has a unique global maximum. For $K = 0$ the landscape has a single maximum, where a maximum is defined as a binary configuration which fitness is larger than the fitness of all its N neighbors (strings that differ from it by a single bit). With increasing K number of such local maxima increases (see Fig. 2), and at extreme value $K = N - 1$ results in a landscape in which the fitness of neighboring points are uncorrelated and the average number of maxima is $2^N/(N + 1)$. As will be demonstrated next, increasing the ruggedness of the landscape (number of local maxima) [39] significantly affects performance of individual search, making it less and less effective as K rises.

Following convention established in earlier papers [3], [19], we normalize the fitness values by the maximum observed value in the landscape, to obtain relative fitness between 0 and 1, with the global maximum set to 1. Next, we rescale these values by raising each number to the power of 8. This monotonic transformation creates larger differences in the upper range of fitness values but leaves other features of the landscape unchanged.

B. Social network structure

Several computational [3], [19], [33] and experimental [25], [38] studies have explored the connection between structure of a team and outcomes of collective problem solving, focusing the attention on networks that facilitate or hinder the exchange of information at the group level. In order to test a wide spectrum of possible structures, we follow the approach of Mason *et al.* [25] and starting from a random regular graph with n nodes and a fixed degree k , we perform degree-preserving rewirings of edges which optimize specific measure of interest.

Since recent studies report contradictory results with regard to efficient (fast information spread) and inefficient (slower information diffusion) structures [3], [19], [25], we select to maximize betweenness centrality as means of creating an inefficient network (INEFF), and minimize average clustering as a way of obtaining an efficient topology (EFF). Additionally, we include in the study a maximally connected network, an all-to-all structure (ATA) and a regular graph (REG), which occupies a middle ground between efficient and inefficient structures, since it is characterized by large clustering and long average paths.

In order to study the connection between size of a social group and its performance in a collective learning task we vary density of these networks by changing the size n of the group,




Figure 2. Landscapes generated by the NK-model vary in ruggedness, as parameter K defining degree of interdependence between N problem variables increases. (a) Low values of K result in a smooth landscape characterized by few local optima, (b) while high K values lead to a landscape with numerous local optima. (c) Agents exploring the landscape in search for global optimum can share information through their social network and avoid mistaking local optimum for the global one. Three dimensional landscapes presented here serve as illustrations of the NK-landscape concept, however real NK-landscape is an N -dimensional construct.

Algorithm 1: Outline of static learning strategies

```
for Time do
  for NodeIndex do
    SampleNeighbors=NNeighborsFitness(Sample)
    if Learning=BEST then
      | TestFitness=max(SampleNeighbors)
    end
    if Learning=FREQ then
      | TestFitness=max(Count(SampleNeighbors))
    end
    if TestFitness > CurrentFitness then
      | CurrentFitness ← TestFitness
    else
      SearchFitness=IndividualSearch
      if SearchFitness > CurrentFitness then
        | CurrentFitness ← SearchFitness
      else
        | KeepState
      end
    end
  end
end
end
```

keeping degree k constant in all cases besides ATA. Figure 4 illustrates basic topological properties of these four networks as a function of their density.

C. Learning rules

In earlier sections we defined two aspects of collaborative problem solving: the structure of networks used to share information between agents and the complexity of problems teams are trying to solve. Here we define learning strategies (see Alg. 1) adopted by teammates in order to localize optimal solutions.

We start from randomly placing n agents of a social network on a given NK-landscape. Next, at each time step of the simulation, all agents synchronously apply a social learning step, followed if necessary by an individual learning step. In the social learning, agents consult a subset of their nearest neighbors, and make decision of either selecting the solution of the best performing neighbor (BEST) [19] or the most frequent solution in the sample (FREQ) [3]. If solutions obtained through social learning are associated with higher fitness value, an agent imitates its better performing neighbor and replaces its solution with the one characterized by higher fitness. However, if the socially acquired solution is no better, the agent switches to individual learning, which consists of examining the fitness value of a solution differing in one randomly chosen bit from the current solution. As in the social learning, the solution obtained through individual learning is accepted only if its fitness is higher than the currently held option. In cases when both social and individual learning fails at providing a better solution, an agent preserves its current configuration.

The main property of learning strategies outlined above is that they are static, performing the same sequence of steps

at the beginning, throughout and at the end of the learning process. Here we introduce two additional learning strategies, which in very basic terms introduce a temporal component to learning. In the first one (EXP), we perform only individual learning steps until at least 90% of agents are unable to locate a better solution than the one currently held. Then all agents switch to social learning by imitating the best among their neighbors. The second strategy is motivated by experimental observations [5] of improved collective learning outcomes when social influence is intermittent (INTER). Thus, in this learning strategy, in an periodic fashion, agents perform a set number of individual learning steps, followed by set number of social learning steps (imitating the best neighbor).

D. Modeling trust

Finally, we consider how different levels of trust between agents within a social network influence the learning strategy and collective problems solving abilities of groups. Trust between agents is represented as a weighted edge applied to a neighbour's fitness value during the social learning phase. This simple representation of trust alludes to the dynamic nature of real world information sharing where the validity and timeliness of information shared can not always be guaranteed [32]. Both the value of the information being shared and the trust in the source of the information are considered in the decision by each agent to leverage social learning or not.

Three levels of average trust within a regular network structure (REG) are explored; HIGH ($\mu = 0.84, \sigma = 0.14$), EVEN ($\mu = 0.5, \sigma = 0.23$) and LOW ($\mu = 0.16, \sigma = 0.14$), where agents consider the best solution from all k neighbors during the social learning phase (BEST) for increasing levels of problem complexity. The edge weights between neighbors are chosen randomly for each level of average trust from three different beta distributions defined on the interval $[0, 1]$ (see Fig. 6 (a-c)). These are compared to COMPLETE ($\mu = 1, \sigma = 0$) average trust in the same network, equivalent to all connected nodes having edge weights of 1. An edge weight of 0 indicates there is no trust between an agent and its neighbor, equivalent to no connection, while an edge weight of 1 indicates a neighbor who is completely trusted.

The amount of trust between two neighbors influences the likelihood that information shared through social learning will be adopted via imitation. Each edge weight is applied to the associated agent's fitness value, effectively downgrading the fitness value for each neighbor being considered during the social learning phase. The level of trust stays constant within the social network for each time step of the numerical simulation.

E. Numerical simulations

For chosen parameters of the social network and the NK-landscape, as well as for given learning strategy, we start a single realization of the collective learning task by assigning random starting solutions to n agents of the social network. Next, at each time step, agents use the strategies described above to learn about the environment. We repeat this process for 200 time steps and we quantify learning by registering

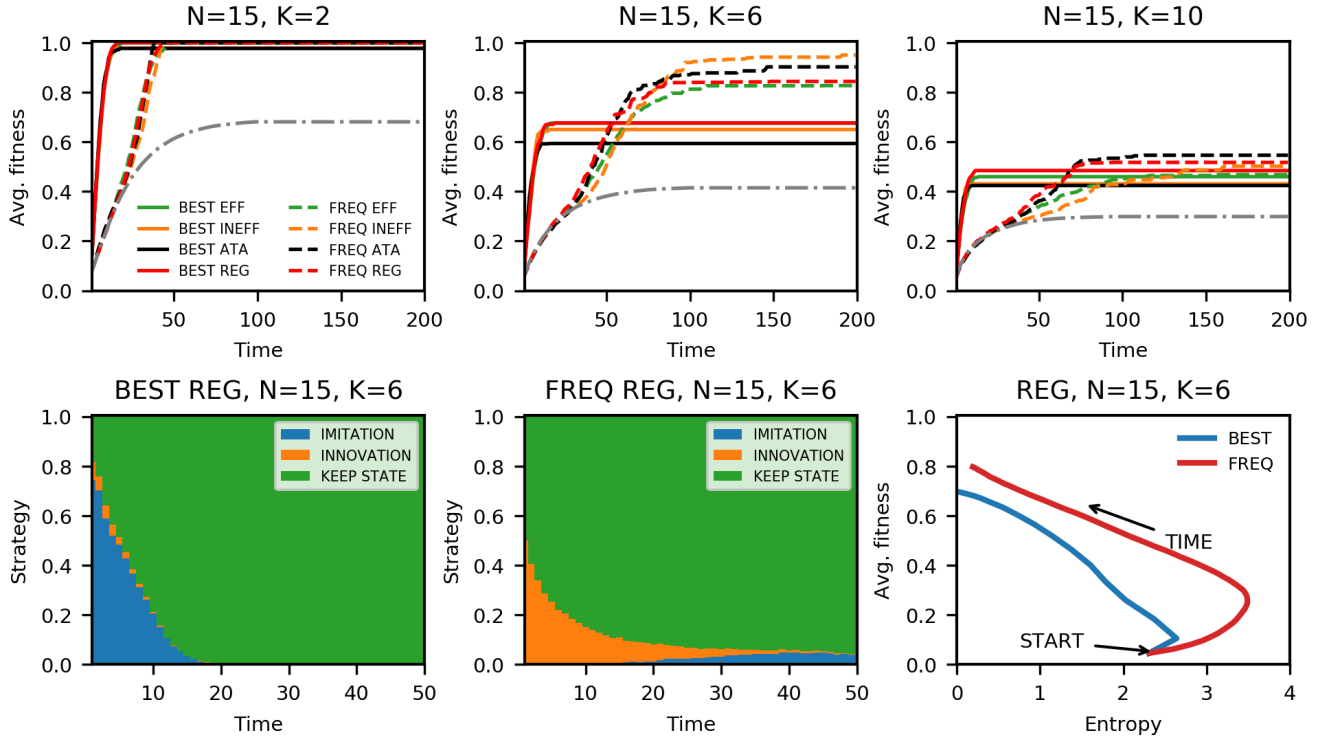


Figure 3. Learning strategy affects success of collective search process by shifting the exploitation-exploration balance. Under BEST strategy agents imitate each other extensively, which leads to a lower performance in high K landscapes. FREQ strategy introduces more innovation in the initial phase of the search process, followed by imitation at later times, which results in better performance at intermediate K values. Differences in social network structure have less pronounced effect on the search process, however performing a social search is always better than searching alone (grey dot-dash line corresponds to the average fitness achieved by agents searching individually). Results presented here are obtained for social networks of $n = 100$ agents, connected with degree $k = 20$ for all networks besides all-to-all (ATA). In all cases comparison within local neighborhood is done using $s = 3$ randomly selected neighbors.

fitness trajectories of all agents, $F_i(t)$, denoting consecutive fitness values adopted by the i -th agent as it explores the landscape. An average over all agents, $F(t) = \frac{1}{n} \sum_{i=1}^n F_i(t)$, describes performance of the entire group, and due to the adopted normalization, $F(t) = 1$ marks the case when all team members found the optimal solution. In order to take into account statistical variation of adopted social networks and NK-landscapes, we repeat the above simulation 100 times. We then use an ensemble average, $\langle F(t) \rangle$, to report results for the given combination of strategy, problem complexity and network structure.

Additionally, in order to quantify the exploration-exploitation patterns created by different learning configurations, we monitor fractions of the agents which at each time step update their fitness value through one of three possible actions: social learning (imitation, exploitation), individual learning (innovation, exploration) or preserving current state. Furthermore, we relate the group performance to observed imitation-innovation ratio by quantifying the amount of information available to the teammates during learning. We adopt Shannon entropy, $H_i(t) = -\sum_j p_j \ln p_j$, to measure the information available to the i -th agent. The p_j is the distribution of fitness values of nearest neighbors of the i -th agent, and similarly to fitness trajectories, we use an ensemble average, $\langle H(t) \rangle$, in further

discussion.

III. RESULTS

A. Exploration-exploitation patterns

Figure 3 shows the average performance achieved by two learning strategies: BEST and FREQ for environments of increasing complexity adopting collective learning in social networks of different structure. In all cases strategies relying on the best member lead to the fastest convergence and consistently collaborative learning outperforms search performed by individuals. Furthermore, the pace of learning based on FREQ strategy is always slower, but in landscapes of medium complexity ($K = 6$) it significantly outperforms other approaches. The success of both learning strategies in simple environments ($K = 2$) stems from the fact that those landscapes contain few local optima, making misidentification of local optima for the global one much less likely. On the other hand, highly complex landscapes ($K = 10$) consist of a significantly larger number of local maxima, resulting in poor group performance, independent from adopted learning strategy.

We have shown that different learning strategies can achieve very different performance within the same environment. At the same time, it is known that different network structures also affect the performance by changing relative use of exploration

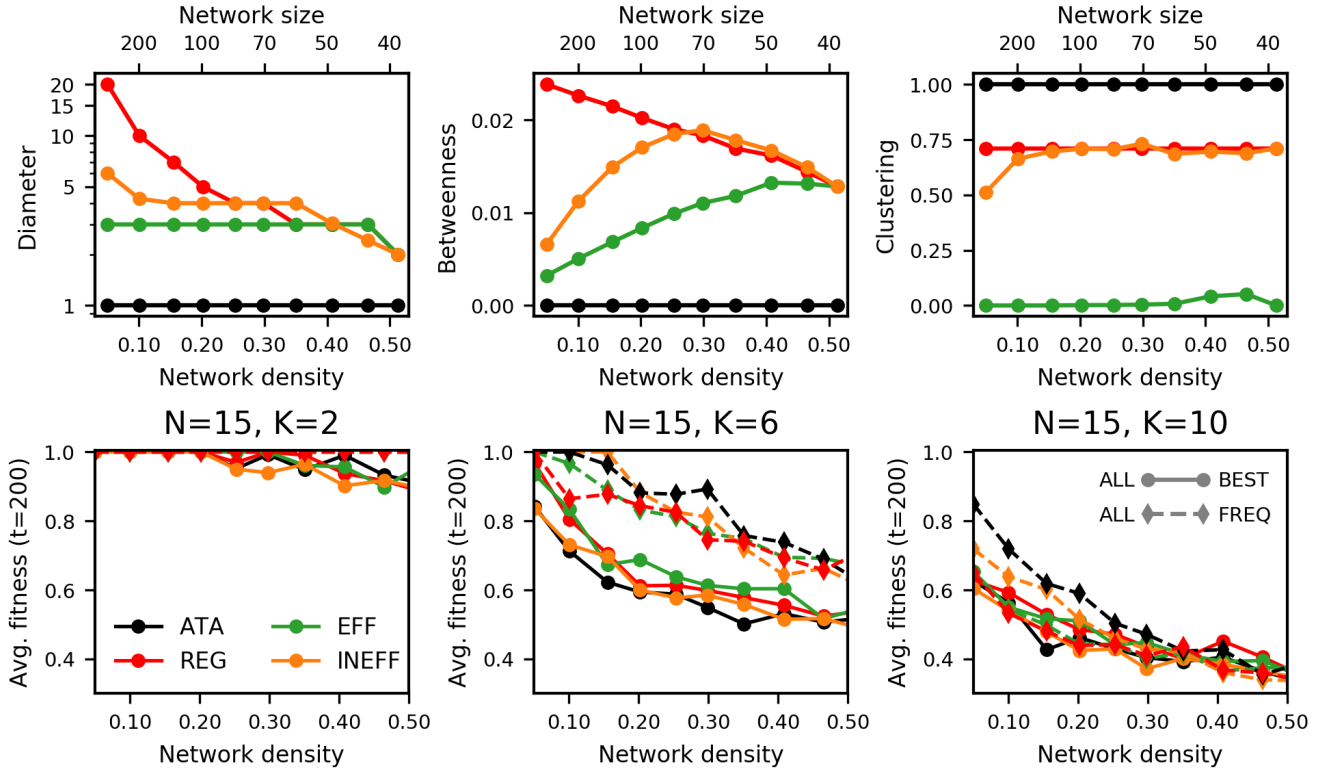


Figure 4. Processes such as diffusion and spreading of information are commonly facilitated by well-connected networks (efficient). However the collective search process benefits primarily from decreasing network density, or more precisely increasing network size. Four analyzed networks differ in their diameter, average clustering and average betweenness centrality, but those differences do not significantly affect the collective search process. Notice: density of the all-to-all network (ATA) is always 1, and the values plotted above correspond to ATA networks of the same number of nodes as other three cases. Degree of networks used here is $k = 20$ for all networks besides ATA.

and exploitation in the group [3], [19]. Here, however, we see that network structure influences group performance in very limited fashion, playing secondary role when contrasted with impact caused by different learning strategies. In particular, the efficiency of topologies, defined in terms of short path lengths, negatively affects learning outcomes.

In order to understand how particular learning strategies influence search outcomes we study by what means (imitation or innovation) agents update their state during the simulation. The fact that the BEST strategy favours imitation, when the FREQ rule promotes innovation is clearly visible on Fig. 3. This different exploration-exploitation pattern is responsible for the success of FREQ over BEST rule. Learning through innovation means exploring new solutions, while learning through imitation corresponds to propagating known answers. The latter does not produce new knowledge and is limited to choosing the best of available solutions, while the former is based on exploration of the environment, increasing chances that over time truly good solution will be found.

We capture this dichotomy by analyzing the relation between information (entropy) available to the agents and average performance. The period of innovation in FREQ strategy corresponds to an increase in entropy of solutions present in the network, and thus in the end is related to better outcomes.

The BEST rule, on the other hand, is characterized mostly by decrease of entropy, directly caused by basing learning on imitation.

B. Team size

In the era of collaboration and teamwork, one key question is what is the optimal team size for achieving the best group performance [9], [11]. In the context of collective learning modeled as search on NK-landscapes, the answer appears to be that large and sparse teams perform the best (see Fig. 4). Similarly to observations reported in other optimization tasks [2], high density of connections is not correlated with improved performance, as one can clearly see by comparing results obtained for a fully connected network (ATA) to ones obtained for sparser topologies. Moreover, the observation that decrease in network density improves search outcomes stems almost entirely from the fact that decreased density signifies an increase in network size, which in turn means that a larger number of agents explore the environment. Additionally, the performance of a team of given size is a function of problem complexity, where highly complex problems require significantly larger teams to perform well. It remains to be seen if the need for increasing team size can be effectively

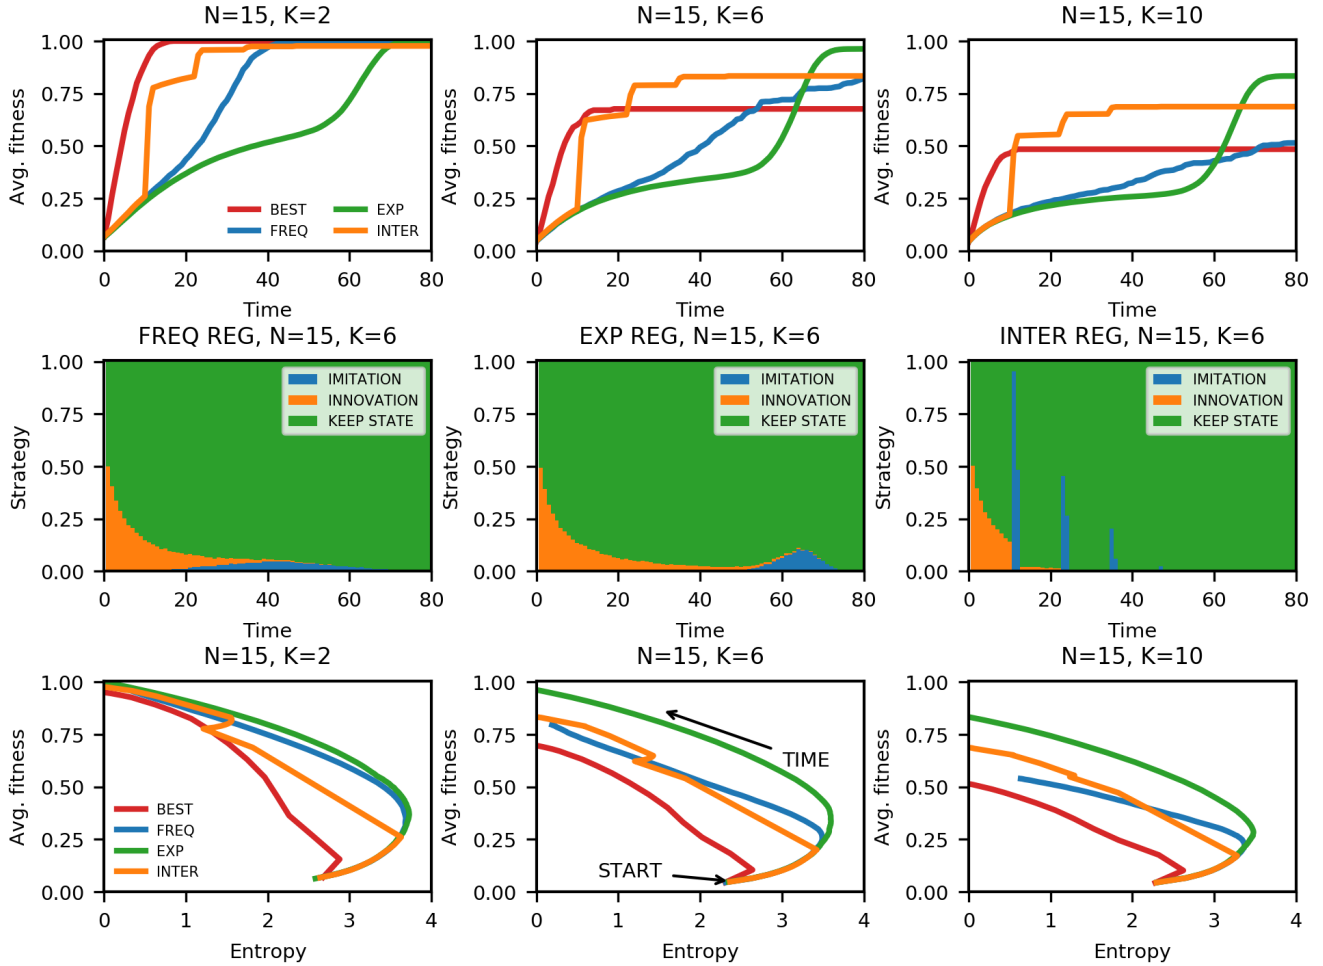


Figure 5. Success of search strategies is strongly linked with the amount of information generated during the search process. Strategies prioritizing innovation (exploration) over imitation (exploitation) in the initial phase of the search (FREQ, EXP) are characterized by an increase in the entropy of fitness values observed by the agents. Imitation strategy acts as a force reducing the information gained by the agents and thus needs to be applied strategically during the search process, as demonstrated e.g. by intermittent (INTER) strategy. In this case imitation is applied sparsely between long intervals of innovation, a procedure which provides good accuracy as well as it improves speed of the search process. Results presented here are obtained for a regular network (REG) of $n = 100$ nodes and degree $k = 20$.

compensated by more efficient learning rules, since the former approach is not realistically possible in entrepreneurial settings.

C. Dynamic learning rules

Observations of animal behavior [14], [26], [29], as well as experiments in humans [5] have demonstrated that dynamic switching between social and individual learning improves performance of the collaborative decision making process. Shifting from group driven activity to individual work is conjectured to act like brainstorming sessions, allowing to identify a wider range of possible solutions, since agents are not influenced by each other at all times.

Figure 5 shows the impact of two simple dynamic learning rules on the performance of the collective learning process. Our results indicate that incorporation of a temporal component into the learning sequence results in comparable or better outcomes than ones obtained in the case of static learning

policies. The EXP rule of executing an extensive individual search followed by social learning operates similarly to FREQ strategy, however due to longer period of individual learning and stronger presence of imitation, it outperforms the latter case. Furthermore, the intermittent strategy is also an improvement over BEST and FREQ rules, in particular since it results in the outcomes equal and better than FREQ rule, but achieves them in a much shorter time.

Finally, we investigate the behavior of entropy generated by those rules. We observe a clear correlation between initial increase in the information (entropy) present in the network and the final outcome of the search process. In all cases a significant adoption of innovation as the learning strategy improves learning outcomes, demonstrating the importance of individual activities in the general scope of a teamwork. Moreover, the entropy-performance relation observed for intermittent strategy suggests that the amount of information necessary

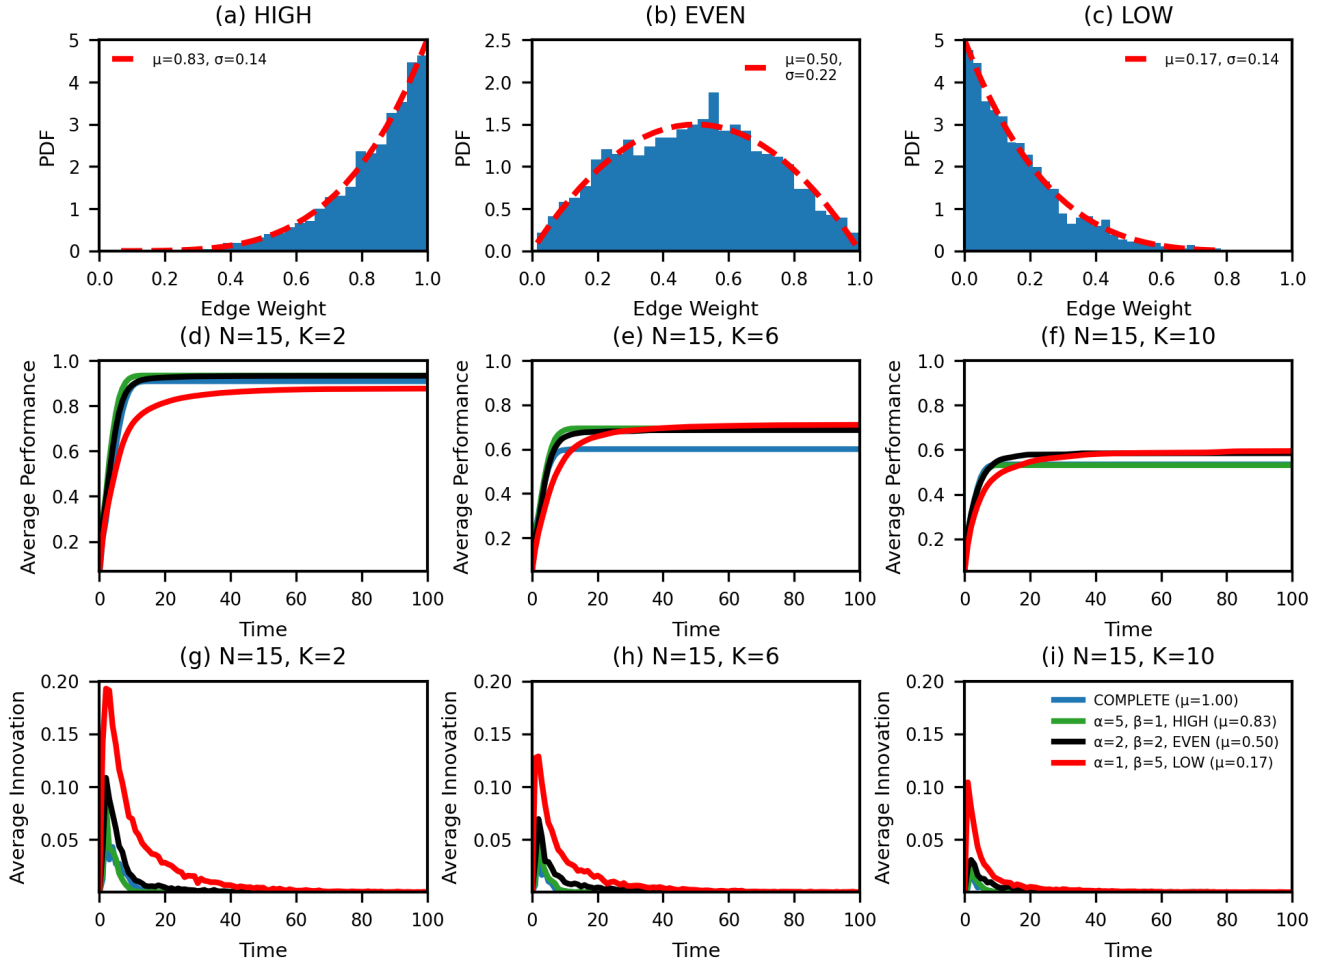


Figure 6. (a-c) Probability density function of edge weights between neighbors in a regular network as sampled from three different beta distribution representing HIGH ($\alpha = 5, \beta = 1$), EVEN ($\alpha = 2, \beta = 2$) and LOW ($\alpha = 1, \beta = 5$) average trust. (d-f) Average fitness and (g-i) average innovation for different levels of trust within a regular network. For simple landscapes ($N = 15, K = 2$), average performance decreases as trust within a network decreases. For complex landscapes ($N = 15, K = 10$), the opposite is observed. In both cases innovation is greater when trust is low within the network as individual search is favored. Results presented here are obtained for a regular network (REG) of $n = 100$ nodes and degree $k = 20$ with all neighbors considered using the BEST social learning strategy.

for identification of an optimal solution is not only a function of the learning rule, but also depends on the complexity of investigated landscape. In this case we observe that comparably good performance is obtained from lower levels of entropy, implying that there exists an optimal amount of information necessary for identification of optimal solutions.

D. How trust influences search process

Figure 6 shows the average performance and average innovation for three levels of average trust within a regular network: HIGH, EVEN and LOW when solving problems of increasing complexity.

Regardless of problem complexity, the learning through innovation increases as the average level of trust in the network moves from HIGH to LOW and the information shared through collective learning is less likely to be trusted. When compared to COMPLETE trust case, this increase in innovation results in

a reduced average performance for simple ($K = 2$) landscapes but an improvement in average performance in complex environments.

The average performance of agents exploring complex landscapes, within a LOW trust network, benefits from greater innovation as the chance of mistaking local optimum for the global one is reduced. For simple landscapes, few local optimum mean avoidance of such mistakes is less of a concern, making imitation a more valuable strategy for increasing average group performance. Within a network of LOW average trust, agents are unlikely to imitate neighbors and instead innovate more. For a simple environment, this ultimately fails to lead to an average performance as strong as when average trust in the network is high.

Social networks with HIGH and EVEN trust marginally outperform networks with COMPLETE trust for simple landscapes due to the effective balance achieved between innovation

and imitation. This suggests that the average trust within a network becomes an influencing factor when the majority of neighbors are not trusted when considering problems in simple environments. In a complex environment, HIGH and COMPLETE trust show very similar performance suggesting that complete trust in all neighbors is not necessary to achieve the same average performance.

In both cases of problem complexity, reducing the average trust in the network, slows down the problem solving capabilities. The exception to this is when average trust is high and the problem environment is simple. In this case, HIGH average trust outperforms COMPLETE trust in both final average performance and number of time steps taken to reach this result.

IV. CONCLUSIONS AND FUTURE WORK

We set out to investigate fundamental aspects of collaborative problem solving by studying how the structure of social networks coupled with different learning rules affects search for solutions to complex problems. We found that learning strategies favoring imitation identify solutions quickly, however their performance decreases with increasing complexity of the problem, severely limiting their usefulness. Moreover, shifting the emphasis of the learning process from diffusion of known solutions to the search for new knowledge results in an improved performance in complex environments, but comes at the cost of slower convergence. Additionally our results indicate that structure of networks used for sharing solutions has limited impact on the group performance. Contrary to results reported in the literature [3], [19], networks characterized by short paths are not significantly more efficient in determining correct solutions than networks defined by long distances. Altogether we find that network structure plays a secondary role in collective learning, being overshadowed by rules of interaction between agents. When considering the effect of trust on the problem solving abilities in a regular network, we observe that lower average trust in the network promotes independent thinking and does not necessarily lead to poor performance.

Furthermore, dynamic interplay between independent exploration and exploitation of social ties proves to be the most beneficial for the collective search process. In particular, we show that exploration-exploitation patterns realized in various learning scenarios can be quantified in terms of information acquired during learning. Those results suggest avenues for further research on efficient collective learning approaches, where search for optimal learning strategies is guided by information-based metrics. For example, this approach could help answer the question of what is the maximum performance of strategies limited to exploitation and exploration, and if additional search mechanisms (e.g radical innovation [4]) are needed to improve learning in complex environments. Another interesting issue is temporal aspect of human interactions, which can affect collective learning outcomes [21]. In this case, an information-based approach could guide search towards

optimal frequency and patterns of interactions [34], extending the issue of an optimal team size into the temporal domain.

Finally, in order to fully answer the question of what types of C2 structures provide most agility, resilience and robustness, we plan to extend our investigations into systems composed of multiple teams, connected in varying forms of hierarchical-like structures. Earlier studies [8], [27], [35] and numerous observations of organizational structures adopted by various institutions [6] suggest that hierarchical networks are the most efficient structures for communication and decision making. However, the work done within the framework of NK-model is limited to two level hierarchies only [35], leaving out more realistic, multilevel networks, which are ubiquitous in military C2 structures. As a continuation of work presented here, we will follow the NATO NEC C2 model [1] in order to study the relation between C2 structures and their efficiency for collaborative problem solving and decision making. We expect that our work will provide insights into designing communication environments required by modern military operations.

ACKNOWLEDGMENT

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] D. ALBERTS, R. HUBER, AND J. MOFFAT, *NATO NEC C2 maturity model*, DOD CCRP, 2010.
- [2] A. ALMAATOUQ, A. NORIEGA-CAMPERO, A. ALOTAIBI, P. M. KRAFFT, M. MOUSSAID, AND A. PENTLAND, *Adaptive social networks promote the wisdom of crowds*, Proceedings of the National Academy of Sciences, 117 (2020), pp. 11379–11386.
- [3] D. BARKOCZI AND M. GALESIC, *Social learning strategies modify the effect of network structure on group performance*, Nature Communications, 7 (2016), p. 13109.
- [4] O. BAUMANN, J. SCHMIDT, AND N. STIEGLITZ, *Effective search in rugged performance landscapes: A review and outlook*, Journal of Management, 45 (2019), pp. 285–318.
- [5] E. BERNSTEIN, J. SHORE, AND D. LAZER, *How intermittent breaks in interaction improve collective intelligence*, Proceedings of the National Academy of Sciences, 115 (2018), pp. 8734–8739.
- [6] S. BILLINGER AND M. WORKIEWICZ, *Fading hierarchies and the emergence of new forms of organization*, Journal of Organization Design, 8 (2019).
- [7] F. A. CSASZAR AND N. SIGGELKOW, *How much to copy? Determinants of effective imitation breadth*, Organization Science, 21 (2010), pp. 661–676.
- [8] P. S. DODDS, D. J. WATTS, AND C. F. SABEL, *Information exchange and the robustness of organizational networks*, Proceedings of the National Academy of Sciences, 100 (2003), pp. 12516–12521.
- [9] R. I. M. DUNBAR, *Coevolution of neocortical size, group size and language in humans*, Behavioral and Brain Sciences, 16 (1993), p. 681–694.
- [10] C. FANG, J. LEE, AND M. A. SCHILLING, *Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning*, Organization Science, 21 (2010), pp. 625–642.

- [11] M. GALESIC, D. BARKOCZI, AND K. KATSIKOPOULOS, *Smaller crowds outperform larger crowds and individuals in realistic task conditions*, *Decision*, 5 (2018), pp. 1–15.
- [12] R. GOLDSTONE AND M. JANSSEN, *Computational models of collective behavior*, *Trends in Cognitive Sciences*, 9 (2005), pp. 424–430.
- [13] C. GOMEZ AND D. LAZER, *Clustering knowledge and dispersing abilities enhances collective problem solving in a network*, *Nature Communications*, 10 (2019), p. 5146.
- [14] T. T. HILLS, P. M. TODD, D. LAZER, A. D. REDISH, AND I. D. COUZIN, *Exploration versus exploitation in space, mind, and society*, *Trends in Cognitive Sciences*, 19 (2015), pp. 46 – 54.
- [15] L. HONG AND S. E. PAGE, *Groups of diverse problem solvers can outperform groups of high-ability problem solvers*, *Proceedings of the National Academy of Sciences*, 101 (2004), pp. 16385–16389.
- [16] B. JAYLES, H.-R. KIM, R. ESCOBEDO, S. CEZERA, A. BLANCHET, T. KAMEDA, C. SIRE, AND G. THERAULAZ, *How social information can improve estimation accuracy in human groups*, *Proceedings of the National Academy of Sciences*, 114 (2017), pp. 12620–12625.
- [17] S. KAUFFMAN AND S. LEVIN, *Towards a general theory of adaptive walks on rugged landscapes*, *Journal of Theoretical Biology*, 128 (1987), pp. 11 – 45.
- [18] S. A. KAUFFMAN AND E. D. WEINBERGER, *The NK model of rugged fitness landscapes and its application to maturation of the immune response*, *Journal of Theoretical Biology*, 141 (1989), pp. 211 – 245.
- [19] D. LAZER AND A. FRIEDMAN, *The network structure of exploration and exploitation*, *Administrative Science Quarterly*, 52 (2007), pp. 667–694.
- [20] D. A. LEVINTHAL, *Adaptation on rugged landscapes*, *Management Science*, 43 (1997), pp. 934–950.
- [21] A. LI, L. ZHOU, Q. SU, S. P. CORNELIUS, Y. Y. LIU, L. WANG, AND S. A. LEVIN, *Evolution of cooperation on temporal networks*, *Nature Communications*, 11 (2020), p. 2259.
- [22] J. LORENZ, H. RAUHUT, F. SCHWEITZER, AND D. HELBING, *How social influence can undermine the wisdom of crowd effect*, *Proceedings of the National Academy of Sciences*, 108 (2011), pp. 9020–9025.
- [23] A. MAO, W. MASON, S. SURI, AND D. J. WATTS, *An experimental study of team size and performance on a complex task*, *PLoS ONE*, 11 (2016).
- [24] J. G. MARCH, *Exploration and exploitation in organizational learning*, *Organization Science*, 2 (1991), pp. 71–87.
- [25] W. MASON AND D. J. WATTS, *Collaborative learning in networks*, *Proceedings of the National Academy of Sciences*, 109 (2012), pp. 764–769.
- [26] K. MEHLHORN, B. R. NEWELL, P. M. TODD, M. D. LEE, K. MORGAN, V. A. BRAITHWAITE, D. HAUSMANN, K. FIEDLER, AND C. GONZALEZ, *Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures*, *Decision*, 2 (2015), pp. 191–215.
- [27] J. MIHM, C. H. LOCH, D. WILKINSON, AND B. A. HUBERMAN, *Hierarchical structure and search in complex organizations*, *Management Science*, 56 (2010), pp. 831–848.
- [28] E. MIU, N. GULLEY, K. LALAND, AND L. RENDELL, *Innovation and cumulative culture through tweaks and leaps in online programming contests*, *Nature Communications*, 1 (2018), pp. 78 – 90.
- [29] L. O’BYRAN, M. BEIER, AND E. SALAS, *How approaches to animal swarm intelligence can improve the study of collective intelligence in human teams*, *Journal of Intelligence*, 8 (2020).
- [30] I. RAHWAN, D. KRASNOSHTAN, A. SHARIFF, AND J.-F. BONNEFON, *Analytical reasoning task reveals limits of social learning in networks*, *Journal of The Royal Society Interface*, 11 (2014), p. 20131211.
- [31] L. RENDELL, R. BOYD, D. COWNDEN, M. ENQUIST, K. ERIKSSON, M. W. FELDMAN, L. FOGARTY, S. GHIRLANDA, T. LILICRAP, AND K. N. LALAND, *Why copy others? Insights from the social learning strategies tournament*, *Science*, 328 (2010), pp. 208–213.
- [32] W. SHERCHAN, S. NEPAL, AND C. PARIS, *A survey of trust in social networks*, *ACM Comput. Surv.*, 45 (2013).
- [33] J. SHORE, E. BERNSTEIN, AND D. LAZER, *Facts and figuring: An experimental investigation of network structure and performance in information and solution spaces*, *Organization Science*, 26 (2015), pp. 1432–1446.
- [34] N. SIGGELKOW AND D. A. LEVINTHAL, *Temporarily divide to conquer: Centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation*, *Organization Science*, 14 (2003), pp. 650–669.
- [35] N. SIGGELKOW AND J. W. RIVKIN, *Speed and search: Designing organizations for turbulence and complexity*, *Organization Science*, 16 (2005), pp. 101–122.
- [36] ———, *When exploration backfires: Unintended consequences of multi-level organizational search*, *Academy of Management Journal*, 49 (2006), pp. 779–795.
- [37] H. A. SIMON, *A behavioral model of rational choice*, *The Quarterly Journal of Economics*, 69 (1955), pp. 99–118.
- [38] N. STEFANOVITCH, A. ALSHAMSI, M. CEBRIAN, AND I. RAHWAN, *Error and attack tolerance of collective problem solving: The DARPA Shredder Challenge*, *EPJ Data Science*, 3 (2014).
- [39] J. VAN CLEVE AND D. B. WEISSMAN, *Measuring ruggedness in fitness landscapes*, *Proceedings of the National Academy of Sciences*, 112 (2015), pp. 7345–7346.
- [40] A. W. WOOLLEY, C. F. CHABRIS, A. PENTLAND, N. HASHMI, AND T. W. MALONE, *Evidence for a collective intelligence factor in the performance of human groups*, *Science*, 330 (2010), pp. 686–688.
- [41] S. WUCHTY, B. F. JONES, AND B. UZZI, *The increasing dominance of teams in production of knowledge*, *Science*, 316 (2007), pp. 1036–1039.