Available online at www.sciencedirect.com

ScienceDirect

Physics of Life Reviews 31 (2019) 320–331

PHYSICS of LIFE reviews

www.elsevier.com/locate/plrev

Review

Agent-based models of collective intelligence

Sandro M. Reia [a], André C. Amado [b], José F. Fontanari [a,*]

[a] Instituto de Física de São Carlos, Universidade de São Paulo, Caixa Postal 369, 13560-970 São Carlos, SP, Brazil
[b] Departamento de Física, Universidade Federal de Pernambuco, 50670-901, Recife, PE, Brazil

Available online 7 January 2019
Communicated by Felix Schoeller

Abstract

Collective or group intelligence is manifested in the fact that a team of cooperating agents can solve problems more efficiently than when those agents work in isolation. Although cooperation is, in general, a successful problem solving strategy, it is not clear whether it merely speeds up the time to find the solution, or whether it alters qualitatively the statistical signature of the search for the solution. Here we review and offer insights on two agent-based models of distributed cooperative problem-solving systems, whose task is to solve a cryptarithmetic puzzle. The first model is the imitative learning search in which the agents exchange information on the quality of their partial solutions to the puzzle and imitate the most successful agent in the group. This scenario predicts a very poor performance in the case imitation is too frequent or the group is too large, a phenomenon akin to Groupthink of social psychology. The second model is the blackboard organization in which agents read and post hints on a public blackboard. This brainstorming scenario performs the best when there is a stringent limit to the amount of information that is exhibited on the board. Both cooperative scenarios produce a substantial speed up of the time to solve the puzzle as compared with the situation where the agents work in isolation. The statistical signature of the search, however, is the same as that of the independent search.

Keywords: Collective intelligence; Imitative learning; Blackboard systems

1. Introduction

The development of software packages to tackle almost any technical problem we can think of, as well as of friendly interfaces between people and computers (see, e.g., [1] in this issue), has contributed to the tremendous increase of the productivity of today's society, thus ratifying the 1960s vision of the computer pioneer Doug Engelbart that the human intellect could be augmented with the aid of computers [2]. In fact, if human intelligence is gauged by the capacity and speed to solve problems, which seems a sensible perspective for the scientific and technological milieus, then it is difficult to disagree with Engelbart's rationale.

However, since the general intelligence of a person – the g factor – has strong genetic (about 50%) and developmental (the remaining 50%) components there is little that adults can do to increase their intelligence (see [3] for a

* Corresponding author.
E-mail address: fontanari@ifsc.usp.br (J.F. Fontanari).

https://doi.org/10.1016/j.plrev.2018.10.004

lucid discussion of these controversial issues), thus limiting the growth of the most important element of the human–computer partnership. Although the natural and widely employed way to circumvent this limitation is to consider group work, only recently psychologists have put forward evidence supporting a general collective intelligence factor – the so-called c factor – that explains the group performance on a variety of tasks [4]. Surprisingly, the c factor does not seem to be strongly correlated with the average or maximum individual intelligence of group members. It is correlated instead with the average social sensitivity of the group members [4]. Therefore, we could, in principle, augment the group intelligence by properly selecting the group composition and internal organization.

Actually, the study of the influence of the organization and, in particular, of intra-group communication patterns (i.e., who can communicate with whom) on the problem-solving performance of groups dates back at least to the 1950s [5,7,6] (see [8–10] for more recent contributions). Understanding this influence is, of course, of immense value because problem solving (e.g., drug design, traffic engineering, software development) by task forces represents a substantial portion of the economy of developed countries nowadays [11].

As argued above, we see organizational design as a by-product of the study of collective intelligence or, more specifically, of distributed cooperative problem-solving systems. The key feature of these systems is that their members exchange information about their progress towards the completion of a goal [12,13]. There are many common-sense assumptions in this field, e.g., that a group of cooperating individuals is more efficient than those same individuals working in isolation or that diversity is always beneficial to the group performance, that were not fully scrutinized through a powerful (in the explanatory sense) analytical tool of physics, namely, the mathematical and computational modeling of complex phenomena using minimal models. A minimal model should exhibit a good balance between simplicity and realism and should be successful in reducing a complex collective phenomenon to its functional essence (see [14,15] for use of this approach to elucidate the physics of mind).

In fact, despite the extensive use of optimization heuristics inspired on cooperative systems, such as the particle swarm optimization algorithm [16] and the adaptive culture heuristic [17,18], to search for optimal or near optimal solutions of combinatorial problems, we know little about the factors that make cooperation effective, as well as about the universal character (if any) of the quantitative improvements that results from it [13]. This is so because those heuristics and the problems they are set to solve are too complex to yield to a first-principle analysis. Here we review recent attempts to study distributed cooperative problem-solving systems using minimal models following the research strategy set forth by Huberman in the late 1980s [12,13,19]. We consider two scenarios of distributed cooperative systems where the goal of the agents is to solve a particular cryptarithmetic puzzle [see eq. 1)], which is an illustrative example of the class of constraint satisfaction problems that bases most studies on problem solving [20].

In the first scenario we endow the agents with the capacity to evaluate the goodness of the partial solutions of their peers and imitate the more successful agent in the group. In the context of collective intelligence or global brains, imitative learning is probably the most important factor as neatly expresses this quote by Bloom "Imitative learning acts like a synapse, allowing information to leap the gap from one creature to another" [21]. Since imitation is central to the remarkable success of our species [22,23] (see [24] for a discussion of this issue), we expect that this scenario may be of relevance to the organization of real-world task groups [25]. In fact, we found that if the agents are too propense to imitate their more successful peers or if the group is too large then the group performance is catastrophic when compared with the baseline situation where the agents work independently of each other [26,27]. This is similar to the classic Groupthink phenomenon of social psychology that occurs when everyone in a group starts thinking alike [28]. Avoiding this sort of harmful effect is the task of organizational designers and we have verified that two rather natural interventions are partially effective to circumvent Groupthink, namely, decreasing the connectivity of the agents so as to delay the propagation of misleading information through the system [29], and allowing diversity in the agents' propensities to imitate their peers [30]. However, these interventions have an unwelcome side effect: the degradation of the optimal performance of the group, which is achieved in the case of fully connected homogeneous agents.

The second scenario of distributed cooperative problem-solving systems that we consider here is the blackboard organization that was introduced in the Artificial Intelligence domain in the 1980s and is now part of the AI problem-solving toolkit [31]. In this organization, the agents read and write hints to a central blackboard that can be accessed by all members of the group. Hence the blackboard scenario describes the common view of task-forces as teams of specialists exchanging ideas on possible approaches to solve a problem and displaying the promising suggestions in a public blackboard. In this scenario, there is no need to assume that the agents are capable of quantifying the goodness of their partial solutions to the cryptarithmetic puzzle as in the imitative learning scenario (see, however, [32]). To our knowledge, this was the first distributed cooperative problem-solving system studied using an agent-based minimal model [13]. Contrary to the claims of that original study, however, we found that the search using the blackboard organization exhibits the same statistical signature of the independent search [32]. Most unexpectedly, we found that limiting the amount of information displayed on the board can markedly boost the performance of the group. This is an original result that we offer in this review paper, which illustrates well the power of minimal models to reveal relevant hidden features of complex systems.

The rest of this short review paper is organized as follows. In Section 2 we present the particular cryptarithmetic problem that we use throughout the paper, define the costs associated to the digit-to-letter assignments that are necessary to implement the imitative learning search strategy, and introduce the definition of hint that is necessary for implementing the blackboard organization. In Section 3 we present our measure of the group performance, which is proportional to the time for an agent in the group to find the solution and hence to halt the search. By dividing this time by the size of the state space of the problem and multiplying it by the number of agents in the group we obtain the computational cost of the search, which is our performance measure. In Section 4 we describe the imitative learning search strategy and discuss its performance on the cryptarithmetic puzzle for the simplest case where all agents interact with each other, i.e., the pattern of communication is a fully connected network. In Section 5 we present the minimal model for the blackboard organization and show that the probability distribution of the computational costs is an exponential distribution as in the case of the independent search. The mean computational cost of the search for the standard blackboard organization, however, is about ten times lower than for the independent search. Also in that section, we introduce the limited space blackboard scenario where the agents have to compete for space in the board. In this case, the computational cost can be reduced by another factor of ten using a judicious choice of the blackboard size. Finally, Section 6 is reserved to our concluding remarks and to advance some avenues for future research.

2. The cryptarithmetic problem

Cryptarithmetic problems such as

$$DONALD + GERALD = ROBERT \tag{1}$$

are constraint satisfaction problems in which the task is to find unique digit-to-letter assignments so that the integer numbers represented by the words add up correctly [33]. In the cryptarithmetic problem (1), there are 10! different digit-to-letter assignments, of which only one is the solution to the problem, namely, $A = 4$, $B = 3$, $D = 5$, $E = 9$, $G = 1$, $L = 8$, $N = 6$, $O = 2$, $R = 7$, $T = 0$ so that $DONALD = 526485$, $GERALD = 197485$ and $ROBERT = 723970$. In this paper we will focus only on the cryptarithmetic problem (1) because its state space is the largest possible for this type of puzzle (a cryptarithmetic puzzle has at most 10 different letters) and because it facilitates the replication of our findings. Use of randomly generated cryptarithmetic puzzles as well as distinct optimization problems, such as finding the global maximum of NK-fitness landscapes [34], has yielded the same (qualitative) results [26,27].

The imitative learning search strategy that will be discussed in Sect. 4 requires that we assign a cost to each digit-to-letter assignment, which is viewed as a measure of the goodness of the answer represented by that assignment. A natural choice for the cost function is [35]

$$c = |ROBERT - (DONALD + GERALD)| . \tag{2}$$

For example, the digit-to-letter assignment $A = 0$, $B = 2$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 3$, $T = 5$ yields $ROBERT = 362435$, $DONALD = 967019$ and $GERALD = 843019$ and the cost assigned to it is $c = 1447603$. We should note that the cost value (2) applies to all digit-to-letter assignments except those for which $R = 0$, $D = 0$ and $G = 0$, which are invalid assignments since they violate the rule of the cryptarithmetic puzzles that an integer number should not have the digit 0 at its leftmost position. Hence for those assignments we fix an arbitrary large cost value, namely, $c = 10^8$, so that now they become valid assignments but have the highest cost among all assignments. If the cost of a digit-to-letter assignment is $c = 0$ then it is the solution to the cryptarithmetic problem.

Although we could easily think up clever alternatives to the cost function (2), we recall that our aim is not to design efficient algorithms to solve cryptarithmetic problems but to explore cooperative strategies that improve the efficiency of group work [12]. In that sense, the chosen problem (1) is quite challenging in that it offers many misleading clues – local minima of the cost (2) and wrong hints – which may lure the search away from the solution. However, as already mentioned, we stress that the main advantage of considering a specific problem is the easy to replicate and verify our claims.

At this stage it is convenient to describe the minimal or elementary move in the state space composed of the 10! possible digit-to-letter assignments. Starting from a particular digit-to-letter assignment, say, $A = 0$, $B = 2$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 3$, $T = 5$ we choose two different letters at random and interchange the digits assigned to them. For example, if we pick letters D and T then the assignment that results from the application of the elementary move is $A = 0$, $B = 2$, $D = 5$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 3$, $T = 9$. Any two valid digit-to-letter assignments that are connected by the elementary move are said to be neighbor assignments. Hence each digit-to-letter assignment in the state space of problem (1) has exactly 45 neighbors. Clearly, the repeated application of our elementary move allows us to explore the entire state space of the cryptarithmetic problem. We can check all assignments and their neighbors to find the number of minima, i.e., those assignments that have a cost (strictly) lower than the cost of their neighbors. We find that the cryptarithmetic problem (1) has 102 minima in total: a single global minimum and 101 local minima. We recall, however, that the existence and characteristics of the local minima are strongly dependent on the choices of the cost function and of the elementary move in the state space.

An important feature of cryptarithmetic puzzles, which makes them a testbed for cooperative strategies, is the existence of hints that may hint on the suitability of a particular digit-to-letter assignment. A hint is a set of letters in a same column that add up correctly modulo 10. For example, considering the third column (from left to right) of the problem (1) we have $B = \epsilon + N + R$ where $\epsilon = 0, 1$ and the sum is done modulo 10. The case $\epsilon = 1$ accounts for the possibility that an 1 is carried from the sum of the letters in the fourth column. Of course, for the rightmost column ($D + D = T$) the only possibility is $\epsilon = 0$. For this column there are 9 different hints: ($D = 1$, $T = 2$), ($D = 2$, $T = 4$), ($D = 3$, $T = 6$), etc. Each one of the columns $\epsilon + L + L = R$, $\epsilon + A + A = E$ and $\epsilon + O + E = O$ has 18 different hints (9 for $\epsilon = 0$ and 9 for $\epsilon = 1$), whereas columns $\epsilon + N + R = B$ and $\epsilon + D + G = R$ have 144 different hints (72 for $\epsilon = 0$ and 72 for $\epsilon = 1$) each. Hence there are a total of 351 distinct hints but only 6 of them yield the solution of the puzzle (1). We note that we could further reduce this number by eliminating hints in the leftmost column such that $\epsilon + D + G > 9$ with the sum now done modulo 1. However, we choose not to implement this rule since what matters is that the total number of hints is much smaller than the size of the state space and that the six correct hints are contemplated in our definition of hint. In the blackboard organization, the communication between the agents is achieved by posting and reading hints in a public blackboard [13] and so there is no need to introduce a cost for each digit-to-letter assignment (see [32] for a scheme where the hints are displayed together with the costs of the agents that posted them).

3. The computational cost

The efficiency of the search for the solution of the cryptarithmetic puzzle is measured by the computational cost that is defined as follows. Let us consider a group composed of M agents so that each agent is represented by a digit-to-letter assignment. The agents explore the state space following a search strategy that specifies the rules for updating their digit-to-letter assignments. Each time a randomly chosen agent updates its digit-to-letter assignment we increment the time t by the quantity $\Delta t = 1/M$, so that during the increment from t to $t + 1$ exactly M, not necessarily distinct, agents are updated. The search ends when one of the agents finds the solution to the puzzle and we denote by t^* the time when this happens. Since we expect that t^* will increase with the size of the state space (10! for our cryptarithmetic puzzle) and that, at least for the independent search, it will decrease with the reciprocal of the number of agents, we define the computational cost C of the search as

$$C = Mt^*/10! \tag{3}$$

so that C is on the order of 1 for the independent search, regardless of the group size. Next we study the statistical properties of the computational cost for two cooperative problem-solving scenarios, namely, the imitative learning search and the blackboard organization.

4. The imitative learning search

This search strategy is based on the presumption that the cost (2) offers a clue on the goodness of the digit-to-letter assignment so that it may be advantageous to copy or imitate agents whose assignments have low cost. More pointedly,

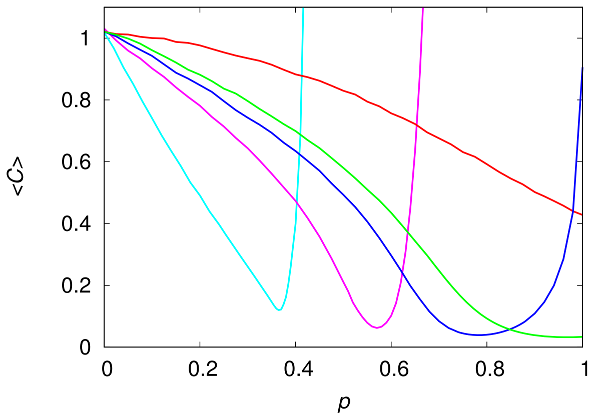

Fig. 1. Mean computational cost $\langle C \rangle$ as function of the imitation probability p for groups of size (top to bottom at $p = 0.2$) $M = 2, 5, 8, 25$ and 200. For the independent search ($p = 0$) we find $\langle C \rangle \approx 1.02$ regardless of the group size.

at time t, a randomly chosen agent – the target agent – can choose between two actions. The first action, which happens with probability $1 - p$, is the elementary or minimal move in the space described in Section 2. The second action, which happens with probability p, is the imitation of the model agent, which is the agent with the lowest cost digit-to-letter assignment in the group at time t. To illustrate the copying process let us assume for the sake of concreteness that the target agent has the assignment $A = 0$, $B = 2$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 3$, $T = 5$ whose cost is $c = 1447603$ and the model agent has the assignment $A = 5$, $B = 3$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 6$, $O = 2$, $R = 7$, $T = 0$ whose cost is $c = 1050568$. In the copying process the target agent selects at random one of the distinct digit-to-letter assignments in the model agent and assimilates it. In our example, the distinct assignments occur for the letters A, B, N, O, R and T. Say that letter B is chosen, so that the target agent has to assimilate the assignment $B = 3$. To do that the target agent simply interchanges the digits assigned to the letters B and R, as in the elementary move, so that the resulting assignment becomes $A = 0$, $B = 3$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 2$, $T = 5$ whose cost is $c = 1545613$. As expected, the result of imitation is the increase of the similarity between the target and the model agents, which may not necessarily lead to a decrease of the cost of the target agent, as in our example. The case $p = 0$ corresponds to the baseline situation where the M agents explore the state space independently.

It is important to note that in the case the target agent is identical to the model agent, and this situation is not uncommon since the imitation process reduces the diversity of the group, the target agent executes the elementary move with probability one. This procedure is different from that used in [26], in which agents identical to the model agent are not updated in the imitation action. Both implementations yield qualitatively similar results, except in the regime where imitation is extremely frequent, i.e., for $p \approx 1$. In particular, for $p = 1$ the implementation in which the model agent is unchanged results in the search being permanently stuck in a local minimum [26], whereas in the implementation in which the model agent executes the elementary move actually leads to the optimal performance for very small groups, as we will show next. The procedure adopted here was used in most studies of the imitative learning search [27,36].

Fig. 1 shows the mean computational cost $\langle C \rangle$ as function of the imitation probability p obtained by averaging over 10^5 independent runs. For groups of size $M < 5$ the performance always improves with increasing p and for those small groups the strategy of always imitating the lowest cost agent (i.e., $p = 1$) is optimal. For $M = 5$ there appears a minimum at $p \approx 0.96$ with computational cost $\langle C \rangle \approx 0.032$ that corresponds to the best performance of the imitative learning search for the entire space of the model parameters M and p. This amounts to more than a thirtyfold improvement on the group performance as compared with the independent search. For large groups, increase of the imitation probability p can lead to catastrophic results due to the trapping of the search around the local minima. This harmful effect appears in large groups only and it is due to the existence of several copies of the model agent carrying a low cost digit-to-letter assignment (local minimum). This makes it very hard to explore other regions of the state space through the elementary move, since the extra copies attract the updated model agent back to the local minimum. However, for small group sizes, the elementary move can easily carry the agents away from the local minima as

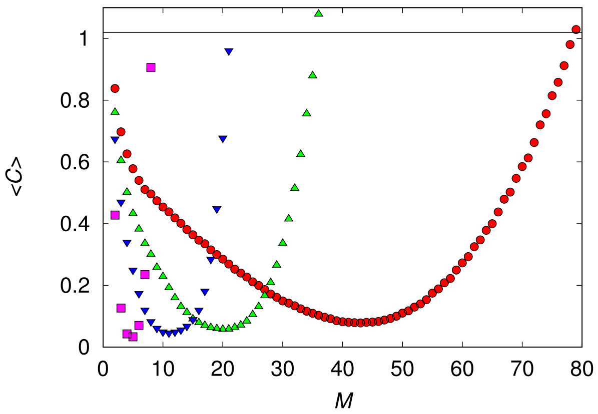

Fig. 2. Mean computational cost $\langle C \rangle$ as function of the group size M for imitation probability $p = 0.5$ (circles), 0.6 (triangles), 0.7 (inverted triangles) and 1 (squares). The horizontal line at $\langle C \rangle = 1.02$ indicates the mean cost for the independent search ($p = 0$).

illustrated in Fig. 1. The optimal performance for $M > 5$, which is determined by the minimum of the curve $\langle C \rangle$ vs. p, degrades smoothly with increasing M. We note that when the model parameters are close to their optimal values, the distribution of probability of the computational costs is well-described by an exponential distribution [26].

Fig. 2 illustrates more neatly the existence of a group size that optimizes the performance of the group for a fixed imitation probability. Hence the conjecture that the efficacy of imitative learning could be a factor determinant of the group size of social animals [26] (see [37,38] for a discussion of the standard selective pressures on group size in nature). As pointed out before, the best performance overall is achieved for small groups with high imitation probability.

The relevant finding revealed in Figs. 1 and 2 is that the group performance can be optimized by tuning the two parameters of the model, namely, the group size M and the imitation probability p. Hence collective intelligence can be augmented through a judicious choice of the behavioral characteristics of the group members, modeled here by the imitation probability p, and of the group size M. As already pointed out, the catastrophic performance exhibited by large groups and high imitation probabilities is akin to the Groupthink phenomenon [28], when everyone in a group starts thinking alike, which can occur when people put unlimited faith in a talented leader (the model agent, in our case). In addition to using small group sizes, there are two other ways to avoid the trapping in the local minima. The first way is decreasing or delaying the influence of the model agent by reducing the connectivity of the network [29]. The second is allowing some diversity in the imitation probabilities of the agents since agents characterized by $p \approx 0$ will rarely be trapped in the local minima [30]. Although these solutions are effective in avoiding the catastrophic performance for large M and p they have the unwanted side effect of degrading the optimal performance, which is obtained using a fully connected influence network with homogeneous imitation probabilities, as in the model described before.

We note that, despite some superficial similarities, the imitative learning search is markedly different from the well-known genetic algorithms [39]. In fact, the elementary move can be seen as the counterpart of mutations with the caveat that in the genetic algorithm mutation is an error of the reproduction process, whereas in the imitative search the elementary move and the imitation procedure are mutually exclusive processes. The analogy between the imitation and the crossover processes is even more far-fetched. The model agent is a mandatory parent in all mates but it contributes a single gene (i.e., a single digit to letter assignment) to the offspring which then replaces the other parent, namely, the target agent. Since the contributed gene is not random – it must be absent in the target agent – the genetic analogy is clearly inappropriate and so the imitative learning search stands on its own as a search strategy.

5. The blackboard organization

The popular view of working groups as teams of specialists that exchange ideas on possible approaches to solve a problem and write the promising lines of investigation in a public display is the inspiration for the blackboard organization [31]. The study of a minimal model of these brainstorming groups, which considers M agents and a central blackboard where the agents can read and write hints, suggested that the blackboard organization could produce a superlinear speedup of the solution time t^* (see Section 3) with respect to the number of group members M [13]. We recall that for the independent search the speedup is linear, i.e., $t^* \propto 1/M$, provided that M is not too large in order to avoid duplication of work. In our problem, duplication of work will occur for unrealistically large groups, $M \gg 10!$, only. If the superlinear speedup claim were correct, then it would offer a nice qualitative evidence of the benefits of cooperation to problem-solving systems. However, recent evidences indicate that whereas the blackboard organization actually produces a significant improvement on the performance of the search, it does not change the nature of the search which exhibits the same characteristics of the independent search and, in particular, the same scaling of t^* with M [32].

An advantage of blackboard systems is that they do not need the introduction of arbitrary cost functions to weight the quality of the digit-to-letter assignments. As it will be clear in our analysis of limited space blackboards, the number of hints exhibited by an assignment is an effective, albeit indirect, measure of its quality. A limited space blackboard can exhibit at most B hints so that, when the blackboard is full, the agents must erase hints to make room for their own hints on the board. Next we describe the dynamics of the blackboard organization.

At the initial time, $t = 1$, all agents' digit-to-letter assignments are selected with equal probability from the pool of the 10! valid assignments. Each agent then checks for all possible hints of its digit-to-letter assignment (see Section 2), singles out the novel hints (i.e., the hints that are not already displayed on the blackboard) and chooses one of them at random to post on the board. In this process, the agent also makes a list of the hints that are displayed on the board and that do not appear in its digit-to-letter assignment. Those are the different hints. Let us assume that the limited space board of size B is full, so the agent must make room to post its selected hint. To do so the agent selects one of the different hints on the board at random and replaces it by its hint. In case the board is not full, the agent simply posts the selected hint on the board. We have tested many variants of this pick-and-replace procedure and found that they produce only negligible quantitative changes on the group performance and so do not affect our conclusions.

Once the initial states of the agents and of the blackboard are set up, the agents can update their digit-to-letter assignments by performing two actions: the elementary move described in Section 2 and the assimilation of one of the hints displayed on the blackboard. The search procedure develops as follows. It begins with a randomly chosen agent – the target agent – picking a hint at random from the blackboard. In the case that there are no hints (i.e., the blackboard is empty), or that the target agent is already using the chosen hint, the agent performs the elementary move; otherwise it assimilates the hint. The assimilation of a hint by the target agent involves the relocation of at most six digits of its digit-to-letter assignment. For example, consider the assimilation of the hint ($N = 1$, $R = 4$, $B = 5$) by an agent that has the assignment $A = 0$, $B = 2$, $D = 9$, $E = 4$, $G = 8$, $L = 1$, $N = 7$, $O = 6$, $R = 3$, $T = 5$. This can be done sequentially using the same assimilation procedure of the imitative learning search described in Section 4. First, the assignment $N = 1$ is assimilated, yielding $A = 0$, $B = 2$, $D = 9$, $E = 4$, $G = 8$, $L = 7$, $N = 1$, $O = 6$, $R = 3$, $T = 5$, then $R = 4$, yielding $A = 0$, $B = 2$, $D = 9$, $E = 3$, $G = 8$, $L = 7$, $N = 1$, $O = 6$, $R = 4$, $T = 5$ and finally $B = 5$ resulting in the digit-to-letter assignment $A = 0$, $B = 5$, $D = 9$, $E = 3$, $G = 8$, $L = 7$, $N = 1$, $O = 6$, $R = 4$, $T = 2$ that exhibits the desired hint. As usual, after the target agent is updated, we increment the time t by the quantity $\Delta t = 1/M$.

After any of the events – elementary move or assimilation of a hint from the blackboard – the target agent checks for all possible hints from its new assignment and executes the pick-and-replace procedure described before. In addition, if the solution of the puzzle is found the search halts and the time $t = t^*$ is recorded.

Fig. 3 shows the distribution of probability $P(C)$ of the computational cost for the independent search and for the blackboard organization in the case there is no space limitation on the blackboard so it can display all 351 distinct hints. Our results show that those distributions are exponential for both search strategies, contrary to the suggestion of [13] that the exponential distribution, which characterizes the independent search, would be replaced by a lognormal distribution for the blackboard organization. From a quantitative perspective, however, the blackboard organization produces a tenfold decrease of the computational cost as compared with the independent search. In particular, $\langle C \rangle \approx 0.10$ for the unlimited space blackboard and $\langle C \rangle \approx 1.02$ for the independent search. We note that the elementary move is slightly less efficient to explore the state space than the replacement of the entire digit-to-letter assignment (global move) used in Refs. [13,32]. This is so because it is not too unlikely to reverse a change made by the elementary move. For example, the probability to reverse a change in a subsequent trial is $2/90$ for the elementary move, whereas it is $1/10!$ for the global move. Interestingly, the replacement of the global by the elementary move has no discernible effect on the performance of the blackboard organization.

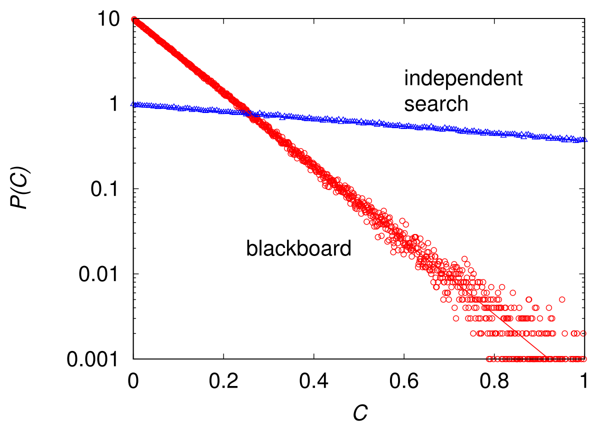

Fig. 3. Probability distribution of the computational cost (3) for the independent search and the blackboard organization with $M = 100$ agents. There is no space limitation on the blackboard, i.e., $B = 351$. These distributions were generated using 10^6 independent runs. The curve fitting the data of the independent search is $P(C) = 0.98\exp(-0.98C)$, whereas the data of the blackboard system is fitted by $P(C) = 10\exp(-10C)$.

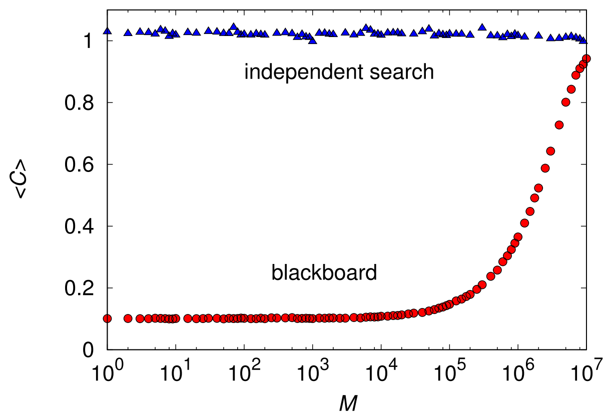

Fig. 4. Mean computational cost $\langle C \rangle$ as function of the system size M for the independent search (triangles) and the blackboard organization (circles). There is no space limitation on the blackboard, i.e., $B = 351$. Each symbol represents the average over 10^5 independent runs. The error bars are smaller than the symbol sizes.

Fig. 4, which shows the mean computational cost as function of the system size, proves that t^* scales with $1/M$ (and hence $\langle C \rangle$ is independent of the system size M) for both the blackboard and the independent search strategies. For the blackboard organization, the increase of the computational cost due to duplication of work occurs for much smaller group sizes than for the independent search since the blackboard reduces the effective size of the state space to be explored by the agents. The main point is that the unlimited blackboard system is not really a cooperative problem-solving system, since once the blackboard is filled out, which happens in a very short time [32], the agents will pick hints on the board and explore the state space independently of each other.

This situation changes dramatically when the size B of the blackboard is limited so the agents have to compete for space to write their hints on the board. At first sight, one would expect that limiting the information available to the agents would harm the group performance. However, Fig. 5 shows that, except for small board sizes ($B < 20$ for the data shown in the figure), increasing the number of hints displayed on the board actually degrades the group performance, regardless of the group size. In addition, for each group size there is a value of B that minimizes the computational cost. In particular, for $M = 1$ we find this optimum at $B = 20$, for $M = 10$ at $B = 7$ and for $M = 100$ at $B = 15$.

To understand the counterintuitive finding that limiting the amount of information displayed on the blackboard improves the group performance, we present in Fig. 6 the probability ϕ that an agent selects a correct hint from the board. This probability is defined as the ratio between the number of correct hints selected and the total number

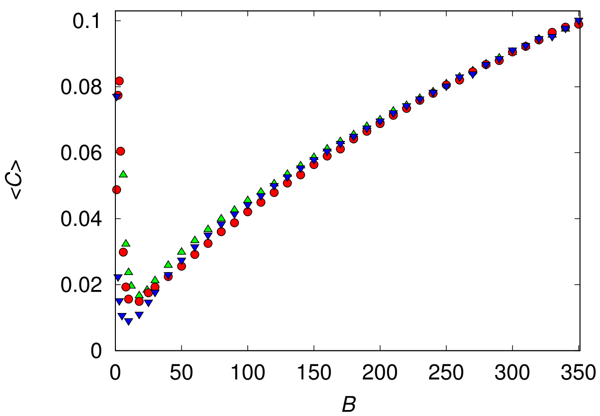

Fig. 5. Mean computational cost $\langle C \rangle$ as function of the blackboard size B for systems of size $M = 1$ (triangles), 10 (inverted triangles) and 100 (circles). Each symbol represents the average over 10^5 independent runs. The error bars are smaller than the symbol sizes.

of hint selections, averaged over all agents during a run. The result is then averaged over 10^5 independent runs. Not surprisingly, Figs. 5 and 6 reveal the strong correlation between $\langle C \rangle$ and ϕ so that the better performance of limited space blackboards is consequence of the higher odds of selecting a correct hint from the board. But the reason these odds are higher for limited boards is not obvious at all. For instance, consider a null model in which the hints displayed on the board of size B are selected randomly without replacement from the pool of 351 hints. Since the cryptarithmetic puzzle (1) has only six correct hints, the probability that the board displays exactly k correct hints is given by an hypergeometric distribution. Now, given that the blackboard displays k correct hints, the probability that the agents selects one of them is simply k/B. Hence the probability that an agent selects a correct hint from a board of size B in this null model is

$$\phi = \sum_{k=0}^{6} \frac{\binom{6}{k}\binom{351-6}{B-k}}{\binom{351}{B}} \frac{k}{B} = \frac{6}{351} \approx 0.017, \tag{4}$$

which does not depend on the board size. We note that in this null model the limitation of the number of hints displayed on the board does not affect the group performance at all, contrary to the naive expectation that it would harm that performance.

The explanation for the dependence on B shown in Fig. 6 (and consequently in Fig. 5) is that the hints displayed on the board are not a random sample of the pool of hints, as assumed in the null model. The somewhat subtle reason for the bias towards the correct hints is that whenever an agent assimilates a hint from the board it must relocate up to six digits of its original digit-to-letter assignment. In doing so, it is very likely to eliminate any previous hints it carried, except if those hints, i.e., the hint copied from the board and the hints that are already part of the agent's assignment, are the correct hints. In that sense, correct hints are insensitive to the radical rearrangement of digits resulting from the assimilation of another correct hint from the board and this explains its higher frequency in the blackboard.

As can be hinted from Fig. 5 the dependence of the computational cost on the number of agents M is quite complex and changes qualitatively for different values of the board size B, so we will leave a detailed discussion of this point to a future contribution. Here we mention only that in a typical situation the computational cost initially decreases with increasing M until it reaches a minimum value, beyond which it begins to increase till it levels off and becomes size-independent, provided M is not too large, as in the case of the unlimited size blackboard (see Fig. 4).

Finally, the fact that the blackboard organization works so well for just a single agent ($M = 1$), which seems to use the blackboard as an external memory to store the hints discovered during its exploration of the state space, indicates that the main role of the blackboard is not the promotion of cooperation between the agents as initially thought [13]: the blackboard serves as a collective memory storage device for the otherwise memoryless agents. We note that for the reputation blackboard, where the hints are posted together with the cost (2) of the agent, the best performance is achieved in the case of a single agent [32]. Interestingly, in the context of the wisdom-of-crowds effect [40] (see also [41]), the performance of a single individual is usually improved if its estimate is taken as the average of its previous estimates – the so-called crowd within [42]. These findings in very distinct contexts make

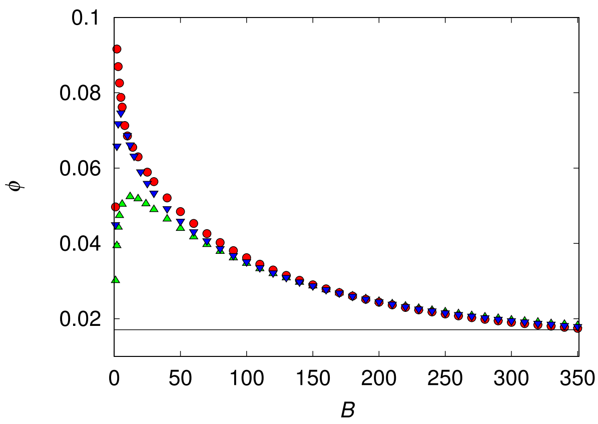

Fig. 6. Probability ϕ that an agent selects a correct hint from the blackboard as function of the blackboard size B for systems of size $M = 1$ (triangles), 10 (inverted triangles) and 100 (circles). Each symbol represents the average over 10^5 independent runs. The error bars are smaller than the symbol sizes. The horizontal line at $\phi = 6/351$ is the prediction of the random blackboard null model.

evident the difficulty to disentangle the memory from the cooperation effects without the aid of the minimal model approach.

6. Conclusion

Rather than advance new search heuristics to solve combinatorial problems, the goal of our approach to study collective intelligence is to assess quantitatively and systematically the potential of cooperation to solve problems in very simplified scenarios. Of course, once the conditions that optimize the efficiency of cooperative work are understood, this knowledge can be used to devise cost-effective search heuristics, which is ultimately the goal of the research on collective intelligence.

Here we have reviewed and offered original insights on two minimal models of distributed cooperative problem-solving systems, namely, the imitative learning search strategy and the blackboard organization. A good criterion to determine whether a mathematical model is minimal or not is the number of model parameters. The imitative learning model has two parameters only: the copy or imitation propensity p of the agents and the group size M. The blackboard organization has two parameters too: the group size M and the blackboard size B. The value of a minimal model (or, for that matter, of any model) should be gauged by the unexpectedness of its predictions. In fact, the optimization of the group performance for a specific group size and the Groupthink-like phenomenon observed in the study of the imitative learning model, as well as the improved performance of the blackboard organization in the case the amount of available information is limited, bear witness to the importance of those models.

From a quantitative perspective, we note that the best performance of the imitative search strategy is achieved for $M = 5$ and $p = 0.96$ and yields the mean computational cost $\langle C \rangle \approx 0.03$ (see Fig. 1), whereas the best performance of the blackboard organization is achieved for $M = B = 10$ with mean cost $\langle C \rangle \approx 0.01$ (see Fig. 5). For the sake of comparison, we recall that the mean computational cost of the independent search is $\langle C \rangle \approx 1.02$, regardless of the group size, so our two cooperative work scenarios produce a substantial boost on the performance of the group of agents.

Although the two minimal models of distributed cooperative problem-solving systems presented here are easy to formulate and simulate in a computer, they exhibit some features that preclude any simple analytical approach as, for instance, the need to select the best digit-to-letter assignment to serve as model for the agents and the non-local effect of assimilating a hint from the blackboard. Perhaps, these obstacles may yield to more powerful and sophisticated mathematical tools, such as the kinetic theory of active particles [43], so as to make the study of collective intelligence more appealing to the mathematics community [44].

To conclude, we note that the key issues that motivated the proposal of the original minimal model of cooperative problem-solving systems remain unanswered [13]. For instance, the question whether cooperative work can alter the statistical signature of the search on the state space of the combinatorial problem is still open, since the computational costs of the models studied here are distributed by exponential probability distributions as in the case of the independent search (see Fig. 3). Moreover, for large group sizes M the time t^* to find the solution decreases with $1/M$ for both the blackboard organization and the independent search (see Fig. 4), whereas it actually increases with increasing M for the imitative search due to the Groupthink phenomenon (see Fig. 2). A qualitative beneficial effect of cooperation should result in the scaling $t^* \propto 1/M^\alpha$ with the exponent $\alpha > 1$, but producing a model with this attribute has proved an elusive task so far.

Acknowledgements

The research of J.F.F. is supported in part by grants 2017/23288-0, Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP) and 305058/2017-7, Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). S.M.R. is supported by FAPESP scholarship 2015/17277-0 and A.C.A. is supported by a CAPES scholarship.

References

[1] Fallani FDV, Bassett DS. Network neuroscience for optimizing brain–computer interfaces. Phys Life Rev 2019;31:304–9. https://doi.org/10.1016/j.plrev.2018.10.001 [in this issue].
[2] Engelbart DC. Augmenting human intellect: a conceptual framework. SRI Summary Report AFOSR-3223. 1962.
[3] Deary IJ. Intelligence: a very short introduction. Oxford, UK: Oxford University Press; 2001.
[4] Woolley AW, Chabris CF, Pentland A, Malone NHT. Evidence for a collective intelligence factor in the performance of human groups. Science 2010;333:686–8.
[5] Bavelas A. Communication patterns in task-oriented groups. J Acoust Soc Am 1950;22:725–30.
[6] Guetzkow H, Simon HA. The impact of certain communication nets upon organization and performance in task-oriented groups. Manag Sci 1955;1:233–50.
[7] Leavitt HJ. Some effects of certain communication patterns on group performance. J Abnorm Soc Psychol 1951;46:38–50.
[8] Mason W, Watts DJ. Collaborative learning in networks. Proc Natl Acad Sci 2012;109:764–9.
[9] Reia SM, Herrmann S, Fontanari JF. Impact of centrality on cooperative processes. Phys Rev E 2017;95:022305.
[10] Reia SM, Fontanari JF. Effect of group organization on the performance of cooperative processes. Ecol Complex 2017;30:47–56.
[11] Page SE. The difference: how the power of diversity creates better groups, firms, schools, and societies. Princeton, NJ: Princeton University Press; 2007.
[12] Huberman BA. The performance of cooperative processes. Physica D 1990;42:38–47.
[13] Clearwater SH, Huberman BA, Hogg T. Cooperative solution of constraint satisfaction problems. Science 1991;254:1181–3.
[14] Perlovsky L. Toward physics of the mind: concepts, emotions, consciousness, and symbols. Phys Life Rev 2006;3:23–55.
[15] Schoeller F, Perlovsky L, Arseniev D. Physics of mind: experimental confirmations of theoretical predictions. Phys Life Rev 2018. https://doi.org/10.1016/j.plrev.2017.11.021.
[16] Bonabeau E, Dorigo M, Theraulaz G. Swarm intelligence: from natural to artificial systems. Oxford, UK: Oxford University Press; 1999.
[17] Kennedy J. Thinking is social: experiments with the adaptive culture model. J Confl Resolut 1998;42:56–76.
[18] Fontanari JF. Social interaction as a heuristic for combinatorial optimization problems. Phys Rev E 2010;82:056118.
[19] Huberman BA, Hogg T. Phase transitions in artificial intelligence systems. Artif Intell 1987;33:155–71.
[20] Newell A, Simon HA. Human problem solving. Englewood Cliffs, NJ: Prentice-Hall; 1972.
[21] Bloom H. Global brain: the evolution of mass mind from the Big Bang to the 21st century. New York: Wiley; 2001.
[22] Blackmore S. The meme machine. Oxford: Oxford University Press; 2000.
[23] Rendell L, Boyd R, Cownden D, Enquist M, Eriksson K, Feldman MW, et al. Why copy others? Insights from the social learning strategies tournament. Science 2010;328:208–13.
[24] Kivinen O, Piiroinen T. Evolutionary understanding of consciousness – a transactional naturalist approach and methodological relationalism. Phys Life Rev 2019;31:32–43 [in this issue].
[25] Lazer D, Friedman A. The network structure of exploration and exploitation. Adm Sci Q 2007;52:667–94.
[26] Fontanari JF. Imitative learning as a connector of collective brains. PLoS ONE 2014;9:e110517.
[27] Fontanari JF. Exploring NK fitness landscapes using imitative learning. Eur Phys J B 2015;88:251.
[28] Janis IL. Groupthink: psychological studies of policy decisions and fiascoes. Boston: Houghton Mifflin; 1982.
[29] Fontanari JF, Rodrigues FA. Influence of network topology on cooperative problem-solving systems. Theory Biosci 2016;135:101–10.
[30] Fontanari JF. When more of the same is better. Europhys Lett 2016;113:28009.
[31] Corkill DD. Blackboard systems. AI Expert 1991;6:40–7.
[32] Fontanari JF. Reputation blackboard systems. Cogn Syst Res 2018;50:29–35.
[33] Averbach B, Chein O. Problem solving through recreational mathematics. San Francisco: Freeman; 1980.
[34] Kauffman S, Levin S. Towards a general theory of adaptive walks on rugged landscapes. J Theor Biol 1987;128:11–45.
[35] Abbasian R, Mazloom M. Solving cryptarithmetic problems using parallel genetic algorithm. In: Proceedings of the second international conference on computer and electrical engineering. Washington, DC: IEEE Computer Society; 2009. p. 308–12.
[36] Fontanari JF. Awareness improves problem-solving performance. Cogn Syst Res 2017;45:52–8.

[37] Wilson ES. Sociobiology. Cambridge, MA: Harvard University Press; 1975.
[38] Dunbar RIM. Neocortex size as a constraint on group size in primates. J Hum Evol 1992;22:469–93.
[39] Goldberg DE. Genetic algorithms in search, optimization, and machine learning. Reading: Addison-Wesley; 1989.
[40] Galton F. Vox populi. Nature 1907;75:450–1.
[41] Surowiecki J. The wisdom of crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations. New York: Random House; 2004.
[42] Vul E, Pashler H. Measuring the crowd within: probabilistic representations within individuals. Psychol Sci 2008;19:645–7.
[43] Bellomo N, Bianca C, Delitala M. Complexity analysis and mathematical tools towards the modelling of living systems. Phys Life Rev 2009;6:144–75.
[44] Burini D, De Lillo S, Gibelli L. Collective learning dynamics modeling based on the kinetic theory of active particles. Phys Life Rev 2016;16:123–39.