



Imitative Learning as a Connector of Collective Brains

José F. Fontanari*

Instituto de Física de São Carlos, Universidade de São Paulo, São Carlos, Brazil

Abstract

The notion that cooperation can aid a group of agents to solve problems more efficiently than if those agents worked in isolation is prevalent in computer science and business circles. Here we consider a primordial form of cooperation – imitative learning – that allows an effective exchange of information between agents, which are viewed as the processing units of a social intelligence system or collective brain. In particular, we use agent-based simulations to study the performance of a group of agents in solving a cryptarithmic problem. An agent can either perform local random moves to explore the solution space of the problem or imitate a model agent – the best performing agent in its influence network. There is a trade-off between the number of agents N and the imitation probability p , and for the optimal balance between these parameters we observe a thirtyfold diminution in the computational cost to find the solution of the cryptarithmic problem as compared with the independent search. If those parameters are chosen far from the optimal setting, however, then imitative learning can impair greatly the performance of the group.

Citation: Fontanari JF (2014) Imitative Learning as a Connector of Collective Brains. PLoS ONE 9(10): e110517. doi:10.1371/journal.pone.0110517

Editor: Long Wang, Peking University, China

Received: August 4, 2014; **Accepted:** September 15, 2014; **Published:** October 16, 2014

Copyright: © 2014 José F. Fontanari. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability: The author confirms that all data underlying the findings are fully available without restriction. All relevant data are within the paper.

Funding: Funding provided by grant 2013/17131-0, São Paulo Research Foundation (FAPESP) and grant 303979/2013-5, Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The author has declared that no competing interests exist.

* Email: fontanari@ifsc.usp.br

Introduction

Imitative learning or, more generally, social learning offers a means whereby information can be transferred between biological or artificial agents, thus being a crucial factor for the emergence of social intelligence or collective brains [1]. Its relevance in this context is neatly expressed by Bloom: “Imitative learning acts like a synapse, allowing information to leap the gap from one creature to another” [2]. Not surprisingly, the advantages of this learning strategy were perceived and exploited by nature well before the advent of the human species as attested by its widespread use in the animal kingdom [3–6]. Regarding human behavior, we note that imitation as a mechanism of social learning was extensively studied by Bandura in the 1960s [7,8] and that the sociocognitive approach to mental processing holds that all mental activity involves either representations of other people or the use of artifacts that have a social history [9,10].

Social learning has inspired the design of several optimization techniques, such as the particle swarm optimization algorithm [11,12] and the adaptive culture heuristic [13,14]. Despite the success of these heuristics in producing optimal or near optimal solutions to combinatorial optimization 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 [15]. The reason is probably that those heuristics and the problems they are set to solve are too complex to yield to a first-principle analysis. In this contribution we address these issues by tackling a simple combinatorial problem and by endowing the agents with straightforward search strategies in which the strength of collaboration is controlled by a single parameter of the model.

The combinatorial problem we consider here is a cryptarithmic puzzle, i.e., a code in which the digits of the integer numbers in a sum are replaced by letters of the alphabet [16,17]. The challenge is to find an assignment between letters and digits that satisfies the constraints of arithmetics as well as the condition that two different letters cannot be assigned to the same digit. In this sense, cryptarithmic puzzles are typical of constraint satisfaction problems which play a central role in our understanding of human and computer problem solving competencies [18,19].

We solve the cryptarithmic problem using a group of N agents which, in addition to the capacity to carry out random local searches, can learn from (or imitate) a model agent – the best performing agent in their influence networks at a given trial. The influence network of each agent is obtained by picking $1 \leq M \leq N-1$ agents at random and without replacement from the $N-1$ remaining agents in the group. The fully connected system corresponds to the case $M=N-1$. The frequency of the imitative or cooperative behavior is determined by the imitation probability parameter $p \geq 0$. Hence our model exhibits two critical ingredients of a collective brain, namely, imitative learning and a dynamic hierarchy among the agents [2].

Our agent-based model conforms to the particle swarm paradigm in that the agents show a tendency to move towards the low cost regions of the solution space which were visited by members of their influence networks [11]. This tendency is a result of the imitation procedure that occurs with probability p . However, we relax the requirement that the agents are more likely to change if the move leads them to a region of lower cost, the so-called Law of Effect [10]. In particular, we allow the agents to move randomly in the solution space with probability $1-p$ and it is this procedure that guarantees that, eventually, one agent will hit the solution of the cryptarithmic problem [15]. We stress

that, whereas the particle swarm algorithm or the adaptive culture heuristic may fail to find the solution of the puzzle because the search can get stuck in a local minimum of the cost landscape, our search procedure, which combines imitation and random changes, always finds the solution. Of course, the issue is how long it takes to do so.

The main performance indicator for the cooperative system is the total number of agent updates necessary to find the solution of the cryptarithmic problem, which we define as the computational cost of the search. The baseline performance corresponds to the case $p=0$ where the N agents explore the solution space independently, resulting in a computational cost that does not depend on the value of N , provided that this value is not too large compared to the size of the solution space. We find that, for a fixed value of the imitation probability p , increasing the number of agents N beyond a certain value impairs the group operation which can then perform much worse than in the case of the independent search. The following of a bad model is the culprit for the poor performance in this case. This harmful effect can be mitigated somewhat by reducing the size M of the influence networks, so as to limit the influence of a model agent to only a fraction of the group. Most significantly, this finding implies that, for fixed p and M , there is a value of group size N that minimizes the computational cost of the search. For instance, in such an optimal setting, say a fully connected system of $N=7$ agents (hence $M=6$) with imitation probability $p=0.6$, we find a thirtyfold decrease of the mean computational cost as compared with the baseline cost.

Methods

First we will describe the particular cryptarithmic problem the agents must solve, explain how the digit-to-letter identifications are encoded in strings and introduce the cost value associated to those strings. We will present also the elementary move that transforms any valid string into an adjacent valid string and so allows the full exploration of the solution space. Once these basic elements are introduced we will describe the mechanism of imitation between agents, thus completing the specification of the agent-based model we use to evaluate the efficacy of imitative learning in solving a complex task.

The cryptarithmic problem

Cryptarithmic problems such as

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

are constraint satisfaction problems in which the task is to find unique digit assignments to each of the letters so that the numbers represented by the words add up correctly [16]. In the cryptarithmic problem (1), there are $10!$ different digit-to-letter assignments, of which only one is the solution to the problem, namely, $A=4, B=3, D=5, E=9, G=1, L=8, N=6, O=2, R=7, T=0$. In fact, with this assignment the cryptarithmic problem (1) is rewritten as the sum $526485 + 197485 = 723970$ which accords with the arithmetic rules. We note that any other one-to-one correspondence between the 10 letters that appear in (1) and the 10 digits would violate those rules. This type of cryptarithmic problem, in which the letters form meaningful words, are also termed alphametics [17] and were popularized in the 1930s by the *Sphinx*, a Belgian journal of recreational mathematics [16]. Of course, from the perspective of evaluating the performance of search heuristics on solving cryptarithmic problems, the meaningfulness of the words is inconsequential, but

in this contribution we will focus mainly on the alphametic problem (1). Nonetheless, we will offer evidence to support the validity of our conclusions by considering a few randomly generated cryptarithmic problems as well.

A non-random search heuristics to solve cryptarithmic problems requires the introduction of some arbitrary quality measure or cost value to each possible digit-to-letter assignment. For the alphametic problem (1) we encode a digit-to-letter assignment by the string $\mathbf{i}=(i_1, i_2, \dots, i_{10})$ where $i_n=0, \dots, 9$ represent the 10 digits and the subscripts $n=1, \dots, 10$ label the letters according to the convention

$$\begin{aligned} 1 &\rightarrow A \\ 2 &\rightarrow B \\ 3 &\rightarrow D \\ 4 &\rightarrow E \\ 5 &\rightarrow G \\ 6 &\rightarrow L \\ 7 &\rightarrow N \\ 8 &\rightarrow O \\ 9 &\rightarrow R \\ 10 &\rightarrow T. \end{aligned} \tag{2}$$

For example, the string $(0,2,9,4,8,1,7,6,3,5)$ corresponds the the digit-to-letter assignment $A=0, B=2, D=9, E=4, G=8, L=1, N=7, O=6, R=3, T=5$. A somewhat natural way to associate a cost to a string \mathbf{i} is through the expression [20]

$$C(\mathbf{i}) = |R - (F + S)| \tag{3}$$

where R is the result of the operation (*ROBERT*), F is the first operand (*DONALD*) and S is the second operand (*GERALD*). In our example we have $R=362435, F=967019$ and $S=843019$ so that the cost associated to string $(0,2,9,4,8,1,7,6,3,5)$ is $C=1447603$. If the cost of a string is $C=0$ then the digit-to-letter assignment coded by that string is the solution of the cryptarithmic problem. We must note that the cost value defined in eq. (3) applies to all strings except those for which $i_3=0$ corresponding to the assignment $D=0, i_5=0$ corresponding to the assignment $G=0$ and $i_9=0$ corresponding to the assignment $R=0$. In principle, those are invalid strings because they violate the rule of the cryptarithmic puzzles that an integer should not have the digit 0 at its leftmost position. For those strings we assign an arbitrary large cost value, namely, $C=10^8$, so that they can be considered valid strings and hence part of the solution space.

In addition to the assignment of the cost values to all $10!$ strings that code the possible digit-to-letter mappings for the alphametic problem (1), we introduce also an elementary move that connects two valid digit-to-letter mappings. We define the elementary move as follows. Starting from a particular digit-to-letter mapping, say $(0,2,9,4,8,1,7,6,3,5)$, we choose two letter labels at random and then interchange the digits assigned to them. For example, say we pick letter labels 1 and 5, then the mapping that results from the application of the elementary move is $(8,2,9,4,0,1,7,6,3,5)$. Clearly, the repeated application of our elementary move is capable of producing all $10!$ strings starting from any valid digit-to-letter mapping.

Imitative learning

The system is composed of N agents or strings which represent valid digit-to-letter assignments as described before. Each agent is connected unidirectionally to exactly $M=1, \dots, N-1$ distinct,

randomly chosen agents in the system. We will refer to those agents as the ‘influencers’ of the target agent. More specifically, for each agent we sample M influencers from the $N-1$ remaining agents without replacement. The extreme case $M=N-1$ corresponds to the fully connected network. An agent has a probability $p \in [0,1)$ of copying a digit-to-letter assignment from a model string in its group of influencers, and probability $1-p$ of performing the elementary move. We choose the model string as the lowest cost string among the M influencers of the target agent. If the cost associated to the target string is lower than the cost of the model string then the copying process is aborted.

To illustrate the copying process let us assume for the sake of concreteness that the target agent is our already familiar example string (0,2,9,4,8,1,7,6,3,5), whose cost is $C=1447603$, and that the model string is (5,3,9,4,8,1,6,2,7,0) 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 string and assimilates it. In our example, the distinct assignments occur at the letter labels $n=1,2,7,8,9,10$. Say that the letter label $n=1$, which corresponds to the assignment $A=5$ according to our convention (2), is chosen. To assimilate this assignment the target agent needs to reassign the digit 0 to the letter label which was previously assigned to digit 5 so that the resulting string becomes (5,2,9,4,8,1,7,6,3,0), whose cost is $C=1448608$. As expected, a result of the imitative learning process is the increase of the similarity between the target and the model strings. The case $p=0$ corresponds to the baseline limit where the N agents perform independent searches. The specific copying procedure proposed here was inspired by the mechanism used to model the influence of an external media [21–23] in Axelrod’s model of culture dissemination [24]. It is important to note that in the case the target string is identical to the model string, as well as in the case the cost of the target string is lower than the cost of the model string, the opportunity of update is wasted.

We may interpret the imitation (or copying) process of a model string as a blackboard cooperation system where a central control exhibits hints (i.e., the lowest cost string) in a public space [15,25], but here we prefer to use the interpretation of learning by imitation in a social context. Nevertheless, since the process of imitation results in an effective collaboration among agents, in the sense that there is an exchange of information between them, we refer to this search strategy as collaborative search to contrast with the independent search which occurs when the copying process is turned off, i.e., the imitation probability p is set to zero.

Search dynamics

We begin by generating the N influence networks, i.e., a group of M influencers for each agent. These networks are kept fixed during the entire search. In this initial stage, at trial number $t=0$, we also associate a random digit-to-letter assignment (a valid string) to each agent and determine its corresponding model string by evaluating and comparing the cost values of its M influencers.

A new trial begins with the choice of the update order of the N agents, so that at the end of the trial all N agents are updated. The agent to be updated – the target agent – has the possibility to imitate its model string or perform the elementary move with probabilities p and $1-p$, respectively. After update, we must re-evaluate the model string status in all groups of influencers to which the target agent belongs. After all N agents are updated we increment the trial number t by one unit and check whether any string has cost zero, in which case the search is halted. The trial number at which the search ends or, alternatively, the number of trial to success is denoted by t^* .

Except for the independent search ($p=0$), the update of the N agents is not strictly a parallel process since the model strings may change several times within a given trial. Nonetheless, since in a single trial all agents are updated, the total number of agent updates at trial t is given by the product Nt .

Results

The efficiency of a search strategy is measured by the total number of agent updates necessary to find the solution of the cryptarithmic problem (i.e., Nt^*) and in the following we will refer to this measure as the computational cost of the search. Since we expect that the typical number of trials to success t^* scales with the size of the solution space (i.e., $10!$), we will present the results in terms of the rescaled variable $\tau = t^*/10!$. For the purpose of comparison we will consider first the independent search strategy where the agents can perform the elementary move only ($p=0$) and then the general cooperative search ($p>0$) where the agents are also allowed to imitate their models.

Independent search

In this case there is no imitation and so the influence networks have no role in the outcome of the search. The main results of the independent search are summarized in Figure 1, which shows the probability distribution $P(N\tau)$ of the rescaled computational cost $N\tau$ of the search for several system sizes. The data is very well fitted by the exponential distribution $P(N\tau) = a \exp(-aN\tau)$ with $a=1/1.14$ which is shown by the solid straight line in the figure.

As expected, the mean rescaled computational cost $\langle N\tau \rangle \approx 1.14$ is insensitive to the system size provided that $N \ll 10!$, but the finding that it does not equal 1 is somewhat surprising. In fact, if we replace our elementary move by a global move in which the entire string is generated randomly at each update then we find that this mean equals 1, as expected. The reason that our elementary move is slightly less efficient than the global move in exploring the solution space is 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 $2/10 \times 1/9 = 2/90$ for the elementary move, whereas it is $1/10!$ for the global move.

Cooperative search

As pointed out before, the cooperation among agents stems from the possibility that they copy potentially relevant digit-to-letter assignments from the model strings in their influence networks. We will consider first the fully connected system where $M=N-1$ and then the partially connected systems where $1 \leq M < N-1$.

Fully connected system. Figure 2 shows how the mean rescaled computational cost is affected by varying the imitation probability p while the number of agents N is kept at a fixed value. For $N=20$ and $p=0.5$ we observe a twentyfold decrease of the mean cost in comparison with the cost of the independent search, which corresponds to $p=0$ and yields $\langle N\tau \rangle \approx 1.14$. This is a remarkable evidence of the power of imitative learning to speed up the search on the solution space of the cryptarithmic problem. In the limit $p \rightarrow 1$ one expects the computational cost to diverge since the solution space cannot be fully explored as the option for the elementary move is never made in this limit. This harmful effect of learning by imitation becomes more pronounced as the number of agents increases.

In the region where the mean computational cost decreases monotonically with increasing p (e.g., $p < 0.5$ for $N=20$) we found that the probability distribution of the computational cost is well




Figure 1. Exponential distribution of the rescaled computational cost for the independent search. Probability distribution $P(N\tau)$ that a search employing N independent agents finds the solution of the cryptarithmic problem (1) using a total of Nt^* updates for $N=5$ (green triangles), $N=10$ (blue inverted triangles) and $N=20$ (red circles). Here $\tau=t^*/10!$ is the ratio between number of trials to success and the size of the solution space. These distributions were generated using 10^5 independent searches for each N . The solid straight line is the exponential distribution $P(N\tau) = a \exp(-aN\tau)$ with $a = 1/1.14$. The influence network size M does not affect these results since imitation is not allowed in this case. doi:10.1371/journal.pone.0110517.g001

described by an exponential distribution, in the sense that the ratio between the standard deviation and the mean is always very close to 1. (We recall that this ratio equals 1 for an exponential distribution.) However, in the region where $\langle N\tau \rangle$ increases with increasing p we found that in the low cost regime $P(N\tau)$ gives values significantly greater than those predicted by an exponential distribution, as illustrated in Figure 3, though those values are not greater than those obtained in the case of the independent search (see Figure 1).

The effect of increasing the number of agents N for a fixed value of the imitation probability p is summarized in Figure 4. The mean computational cost of the cooperative system exhibits a non-monotonic dependence on N , except in the case of the independent search ($p=0$) when it takes on a constant value. The benefit of cooperation is seen in this figure by the initial decrease of the computational cost as the number of agents increases. However, for all $p>0$ we find that the presence of too many agents greatly harms the performance of the system and that for a fixed $p>0$ there exists an optimum value of N that maximizes the search efficiency of the cooperative system. For instance, although not shown in the scale of Figure 4, the minimum computational cost for $p=0.3$ occurs at $N \approx 270$. The efficiency at this optimum, however, is not affected significantly by the choice of the parameters N and p . In other words, the costs corresponding to the minima shown in Figures 2 and 4 are not very sensitive to changes in N and p , respectively. In particular, for the parameter settings we have explored, the best efficiency $\langle N\tau \rangle \approx 0.041$ is achieved for $N=7$ and $p=0.6$ and amounts to a thirtyfold speed up with respect to the independent search.

We conjecture that the reason the efficiency of the cooperative system deteriorates as N increases beyond its optimum value (e.g., in the range $N>7$ for $p=0.6$ as shown in Figure 4) is that for N not too small there is a good chance that the cost of one of the strings is significantly lower than the cost of the other $N-1$ strings. Provided p is not too small too, this string may remain as the model string for a few trials thus biasing the search to the vicinity of the model string. In the (typical) case that the model string is far




Figure 2. The effect of the imitation probability on the computational cost of the fully connected system. The symbols represent the mean rescaled computational cost $\langle N\tau \rangle$ for cooperative systems of size $N=20$ (red circles), $N=5$ (green triangles), $N=3$ (blue inverted triangles) and $N=2$ (magenta squares). The independent variable p is the probability that an agent will copy a digit-to-letter assignment from the model string, chosen as the lowest cost string in the entire system. The influence network size is $M = N - 1$. Each symbol represents an average over 10^5 searches and the lines are guides to the eye. The error bars are smaller than the size of the symbols. doi:10.1371/journal.pone.0110517.g002

from the solution of the cryptarithmic problem, imitative learning may lead to the observed impairment of the performance of the cooperative system. In sum, the following of a bad leader is likely the culprit of the poor performance of the system.

To check the validity of this conjecture we calculate the mean number of consecutive trials for which a cost value stays as the lowest cost among the N strings. The procedure to obtain this quantity, which we denote by ζ , is straightforward. At trial $t=0$ we evaluate the cost of the N strings and record the minimal cost among them. Then at the next trial $t=1$, after the N strings are updated, we re-evaluate again their costs and record the minimal cost. If the minimal cost at $t=1$ is different, i.e., greater or less, than the minimal cost at $t=0$ we say that a change event has occurred. The comparison of the values of the minimal costs at consecutive trials is repeated and the cumulative number of change events is recorded until the solution is found at $t=t^*$. The desired quantity ζ is given simply by the ratio between the total number of change events and the total number of trials t^* . Hence for each search we obtain a single value for ζ , which can then be interpreted as the mean number of trials between consecutive change events or as the mean duration of the stases for that search.

In Figure 5 we present the probability distribution $Q(\zeta)$ using 10^5 searches for the imitation probability $p=0.6$ and two representative values of N . Figure 5A shows this distribution for $N=6$, which corresponds to a regime of low computational cost according to Figure 4. We observe a pronounced maximum at $\zeta \approx 3.7$ so that in most searches the model cost remains unaltered for 3 to 5 trials. This is an optimum scenario since no string stays on the top tier long enough to influence the entire system. For $N<6$, we find that $Q(\zeta)$ exhibits a similar shape but the maximum becomes sharper and its location is shifted towards lower values of ζ as N decreases. Figure 5B, which shows the results for $N=17$, reveals a very different scenario: the distribution $Q(\zeta)$ exhibits a plateau indicating that the model cost remains unchanged for hundreds to a few thousands trials. For very large values of ζ , the distribution $Q(\zeta)$ seems to exhibit an exponential




Figure 3. Deviation from the exponential distribution for a large imitation probability. Probability distribution $P(N\tau)$ of the rescaled computational cost for a search employing $N=20$ fully connected agents with imitation probability $p=0.6$. The mean of this distribution is $\langle N\tau \rangle \approx 8.0$. The solid straight line is the fitting function $a \exp(-bN\tau)$ with $a=0.03$ and $b=1/15$ in the regime of large cost. The influence network size is $M=N-1$. The distribution was generated using 10^5 independent searches.
doi:10.1371/journal.pone.0110517.g003

decay to zero, namely, $Q(\zeta) \sim \exp(-0.003\zeta)$. We stress that for the two cases exhibited in Figure 5 the probability that an agent will imitate the model rather than perform an elementary move is the same, namely $p=0.6$, and so the qualitative differences reported in the figure are due solely to the change on the number of agents.

Partially connected system. If the poor performance of large collaborative systems based on imitative learning is due to the influence of bad models then a natural way to reduce this harmful effect is to limit the influence of those models. This was the motivation to introduce the influence networks scheme where each agent picks its model among M randomly chosen agents predetermined at the beginning of the search. In fact, Figure 6 shows that the reduction of the connectivity of the agents increases somewhat the range of values of the imitation probability p for which the cooperative system outperforms the system composed of independent agents. More pointedly, for $N=20$ this range is extended from $p \approx 0.57$ for $M=19$ to $p \approx 0.87$ for $M=1$. In addition, the value of the optimal mean computational cost does not seem to vary significantly with M . Figure 7 offers another perspective on the role of the number of influencers M . It shows that for small values of the imitation probability the fully connected system (i.e., $M=N-1$) exhibits the best performance. However, as p increases (e.g., $p > 0.4$ for $N=40$), the optimal performance is obtained with partially connected systems. Moreover, we found that for any fixed value of $p > 0$ and M the performance of the system is always impaired when the number of agents N is very large. Finally, we note that similarly to our findings for the fully connected system, the probability distribution of the computational cost $P(N\tau)$ departs significantly from an exponential distribution only in the regions where the mean computational cost becomes an increasing function of the control parameters of the model.

Random cryptarithmic problems. In order to verify the generality of our findings, which were obtained for the specific alphabetic problem $DONALD + GERALD = ROBERT$, we have considered a variety of random cryptarithmic problems with 10 letters and a unique solution, so that the sizes of their solution spaces are the same as that of the alphabetic problem.




Figure 4. The effect of group size on the computational cost of the fully connected system. The symbols represent the mean rescaled computational cost $\langle N\tau \rangle$ for the imitation probability $p=0$ (magenta circles), $p=0.3$ (red diamonds), $p=0.4$ (green squares), $p=0.5$ (blue inverted triangles) and $p=0.6$ (cyan triangles). The independent variable N is the number of agents in the system. The influence network size is $M=N-1$. Each symbol represents an average over 10^5 searches and the lines are guides to the eye. The error bars are smaller than the size of the symbols.
doi:10.1371/journal.pone.0110517.g004

The comparison between the mean computational costs to solve four such random problems and our alphabetic problem is shown in Figure 8 for the fully connected system. The results are qualitatively the same, as expected. The alphabetic problem, however, was somewhat easier to solve by the cooperative system than the random problems, perhaps because of the coincidence of the last three letters (“ALD”) in the first and second operands. Interestingly, the independent system ($p=0$) cannot distinguish between the problems but the cooperative system ($p > 0$) can, and this distinction is most pronounced when the parameters are set so as to achieve the optimal performance. It is as if the cooperative system had adapted to the specific task posed to it. We expect that our conclusions remain valid, in a qualitative sense of course, for any constraint satisfaction problem characterized by a rugged cost landscape.

Discussion

Rather than offer any novel method to solve cryptarithmic problems, our aim in this contribution is to assess quantitatively the potential of imitative learning as the underlying mechanism – the critical connector – of collective brains [2]. Here imitative learning is implemented by allowing an agent to copy clues from the best performing agent – the model agent – in its group of influencers. More pointedly, at trial t each agent has the probability p of imitating the model and the probability $1-p$ of executing a random rearrangement of the digit-to-letter mapping which is its guess to the solution of the cryptarithmic problem. In an optimal setting, say a fully connected system of $N=7$ agents with imitation probability $p=0.6$, we find a thirtyfold decrease of the mean number of trials needed to find the solution of the problem (i.e., of the mean computational cost), as compared with the case $p=0$ when the agents search the solution space independently (see Figure 4).

In the optimal setting, as well as in the regions where the computational cost is a decreasing function of the control parameters of the model, the probability distribution of the computational cost is given by an exponential distribution, rather




Figure 5. Probability distribution of the mean duration of the stases in a search. Probability distribution of the mean number of trials ζ for which a cost value stays as the lowest cost among the N solutions in 10^5 searches for the imitation probability $p=0.6$ in a fully connected system. Panel **A**: $N=6$ (low computational cost regime). Panel **B**: $N=17$ (high computational cost regime). The slope of the straight line shown in the semi-log scale of panel **B** is 0.003. The influence network size is $M=N-1$.
doi:10.1371/journal.pone.0110517.g005

than by a lognormal distribution as predicted by a general theory of cooperative processes [15,26]. In fact, the reason the cooperative scheme implemented in [15] is so efficient is that all discovered digit-to-letter assignments that add up correctly modulo 10 for at least one column are permanently exposed as hints in a blackboard for use by all agents, which can pick a hint at each trial. There is no place for any kind of learning in that scenario since in the case there are no hints in the blackboard or the agent has already used the chosen one, the target agent selects at random a complete digit-to-letter mapping, which is totally uncorrelated to its previous mapping. The imitative learning interpretation of our




Figure 6. The effect of the imitation probability on the computational cost of partially connected systems. The symbols represent the mean rescaled computational cost $\langle N\tau \rangle$ for a system composed of $N=20$ agents, each one connected to $M=19$ (red circles), $M=9$ (green triangles), $M=4$ (blue inverted triangles) and $M=1$ (magenta squares) influencers. The independent variable p is the imitation probability. Each symbol represents an average over 10^5 searches and the lines are guides to the eye. The error bars are smaller than the size of the symbols.
doi:10.1371/journal.pone.0110517.g006




Figure 7. The effect of the number of influencers on the mean computational cost. The symbols represent the mean rescaled computational cost $\langle N\tau \rangle$ for a system composed of $N=40$ agents and imitation probability $p=0.5$ (red circles), $p=0.45$ (green triangles), $p=0.4$ (blue inverted triangles) and $p=0.3$ (magenta squares). The independent variable M is the size of the group of influencers of each agent. Each symbol represents an average over 10^5 searches and the lines are guides to the eye. The error bars are smaller than the size of the symbols.
doi:10.1371/journal.pone.0110517.g007

cooperative scheme is only possible because in our case the elementary random move is local (i.e., solely two digit-to-letter assignments are changed in the entire mapping) and therefore preserves the identity of the target agent.

Most significantly, for fixed values of the imitation probability p and of the number of influencers M , we find that increasing the number of agents N beyond a certain quantity impairs the working of the cooperative system, which then performs much worse than if the agents had executed independent searches. Our analysis indicates that the following of a bad model is the culprit of




Figure 8. Computational cost of the alphametic problem and of four random cryptarithmic problems. The mean rescaled computational cost for the alphametic problem $DONALD+GERALD=ROBERT$ (red circles) and for four ten-letter random cryptarithmic problems with a unique solution (blue inverted triangles, magenta squares, cyan diamonds and green triangles). The symbols represent the mean rescaled computational cost $\langle N\tau \rangle$ for a system composed of $N=20$ fully connected agents. The independent variable p is the imitation probability. The influence network size is $M=N-1$. Each symbol represents an average over 10^5 searches and the lines are guides to the eye. The error bars are smaller than the size of the symbols.
doi:10.1371/journal.pone.0110517.g008

the poor performance of the system in this case. In that sense, the efficacy of imitative learning could be a factor determinant of group size [27]. In contrast to the cognitive load that constrains the number of individuals with whom it is possible to maintain stable relationships and leads to Dunbar's number for primates [28], the group size here (i.e., the value of N corresponding to the lowest computational cost) does not stem from a limitation of the neocortical processing capacity of the individuals. Rather, it is a property of the group of agents as a whole, since for any fixed non-vanishing value of the imitation probability, which may be seen as an individual trait, a too large number of agents, which is a group property, will impair the performance of the cooperative system. Of course, if p were allowed to decrease with increasing N then the system could be maintained at the highest level of perform regardless of the group size (see Figures 2 and 4). In other words, in order to perform at the optimal level a system based on imitative learning should decrease the frequency of the interactions among individuals as its size increases.

References

- Nehaniv CL, Dautenhah K (2007) Introduction: the constructive interdisciplinary viewpoint for understanding mechanisms and models of imitation and social learning. In: Nehaniv CL, Dautenhah K, editors, *Imitation and Social Learning in Robots, Humans and Animals*. Cambridge, UK: Cambridge University Press, pp. 1–18.
- Bloom H (2001) *Global Brain: The Evolution of Mass Mind from the Big Bang to the 21st Century*. New York: Wiley.
- Moore BR (1996) The evolution of imitative learning. In: Heyes CM, Galef BG Jr, editors, *Social learning in animals: The roots of culture*. San Diego, CA: Academic Press, pp. 245–265.
- Fiorito G, Scotto P (1992) Observational learning in octopus vulgaris. *Science* 256: 545–547.
- Laland KN, Atton N, Webster MM (2011) From fish to fashion: experimental and theoretical insights into the evolution of culture. *Phil Trans R Soc B* 366: 958–968.
- Heyes C (2012) Grist and mills: on the cultural origins of cultural learning. *Phil Trans R Soc B* 367: 2181–2191.
- Bandura A (1962) Social learning through imitation. In: Jones MR, editor, *Nebraska Symposium on Motivation*. Lincoln, NE: University of Nebraska Press, pp. 211–274.
- Bandura A (1977) *Social learning theory*. New York: Prentice Hall.
- Levine JM, Resnick LB, Higgins ET (1993) Social foundations of cognition. *Annu Rev Psychol* 44: 585–612.
- Kennedy J (1999) Minds and cultures: Particle swarm implications for beings in sociocognitive space. *Adapt Behav* 7: 269–288.
- Kennedy J, Eberhart R (1995) Particle swarm optimization. In: *Proceedings of the IEEE International Conference on Neural Networks*. Washington, DC: IEEE Computer Society, volume 4, pp. 1942–1948.
- Bonabeau E, Dorigo M, Theraulaz G (1999) *Swarm Intelligence: From Natural to Artificial Systems*. Oxford, UK: Oxford University Press.
- Kennedy J (1998) Thinking is social: Experiments with the adaptive culture model. *J Conflict Res* 42: 56–76.
- Fontanari JF (2010) Social interaction as a heuristic for combinatorial optimization problems. *Phys Rev E* 82: 056118.
- Clearwater SH, Huberman BA, Hogg T (1991) Cooperative solution of constraint satisfaction problems. *Science* 254: 1181–1183.
- Averbach B, Chein O (1980) *Problem Solving Through Recreational Mathematics*. San Francisco: Freeman.
- Hunter JA (1976) *Mathematical Brain Teasers*. New York: Dover.
- Newell A, Simon HA (1972) *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Kirkpatrick S, Selman B (1994) Critical behavior in the satisfiability of random boolean expressions. *Science* 264: 1297–1301.
- Abbasian R, Mazloom M (2009) Solving cryptarithmic problems using parallel genetic algorithm. In: *Proceedings of the Second International Conference on Computer and Electrical Engineering*. Washington, DC: IEEE Computer Society, volume 2, pp. 308–312.
- Shibanai Y, Yasuno S, Ishiguro I (2001) Effects of global information feedback on diversity. *J Conflict Res* 45: 80–96.
- González-Avella J, Cosenza M, Eguiluz VM, Miguel MS (2010) Spontaneous ordering against an external field in non-equilibrium systems. *New J Phys* 12: 013010.
- Peres LR, Fontanari JF (2011) The media effect in Axelrod's model explained. *Europhys Lett* 96: 38004.
- Axelrod R (1997) The dissemination of culture: A model with local convergence and global polarization. *J Conflict Res* 41: 203–226.
- Englemore R, Morgan T (1988) *Blackboard Systems*. New York: Addison-Wesley.
- Huberman BA (1990) The performance of cooperative processes. *Physica D* 42: 38–47.
- Wilson E (1975) *Sociobiology*. Cambridge, MA: Harvard University Press.
- Dunbar RIM (1992) Neocortex size as a constraint on group size in primates. *J Human Evol* 22: 469–493.
- Laland KN, Williams K (1998) Social transmission of maladaptive information in the guppy. *Behav Ecol* 9: 493–499.

Author Contributions

Conceived and designed the experiments: JFF. Performed the experiments: JFF. Analyzed the data: JFF. Contributed reagents/materials/analysis tools: JFF. Wrote the paper: JFF.