Factors of Collective Intelligence: How Smart Are Agent Collectives?

Nader Chmait and **David L. Dowe** and **Yuan-Fang Li** and **David G. Green**[1] and **Javier Insa-Cabrera**[2]

Abstract. The dynamics and characteristics behind intelligent cognitive systems lie at the heart of understanding, and devising, successful solutions to a variety of multiagent problems. Despite the extant literature on collective intelligence, important questions like "how does the effectiveness of a collective compare to its isolated members?" and "are there some general rules or properties shaping the spread of intelligence across various cognitive systems and environments?" remain somewhat of a mystery. In this paper we develop the idea of collective intelligence by giving some insight into a range of factors hindering and influencing the effectiveness of interactive cognitive systems. We measure the influence of each examined factor on intelligence independently, and empirically show that collective intelligence is a function of them all simultaneously. We further investigate how the organisational structure of equally sized groups shapes their effectiveness. The outcome is fundamental to the understanding and prediction of the collective performance of multiagent systems, and for quantifying the emergence of intelligence over different environmental settings.

1 INTRODUCTION

Collective intelligence emerges in all sorts of cognitive systems, from natural (e.g., animal and human) to artificial (e.g., software agents and robotics), by cause of diverse social organizations (human societies, efficient markets, social insect colonies, group collaborations via the web, etc.). It seems that the complex structure and operation of these systems hinder our understanding of the dynamics and characteristics behind intelligent collectives, which are fundamental for devising successful models and solutions to a variety of multiagent problems. Despite the extant literature on collective intelligence (CI), the questions consisting of, "how does the effectiveness of a collective compare to its isolated members?" and, more importantly, "are there some general rules shaping the spread of intelligence which can be perceived across different cognitive systems and environments?" remain somewhat of a mystery. Now imagine we had a series of performance tests over which we can administer any type of cognitive system, could we then disclose any patterns or factors at all, explaining the emergence of intelligence among all of these systems? In this paper, we give insight into the main components and characteristics of collective intelligence, by applying formal tests for the purpose of measuring and quantifying the influence of several factors on the collective behaviour and the accuracy of a group of agents, and analysing how the results compare to individual agent scenarios. We attempt to uncover some of the dynamics and circumstances behind intelligent collectives in general, hoping this would reinforce the understanding and prediction of the behaviour of groups, by bringing some new results into the AI community.

[1] Faculty of IT, Clayton, Monash University, Australia, email: {nader.chmait,david.dowe,yuanfang.li,david.green}@monash.edu
[2] Universitat Politècnica de València, Spain, email: jinsa@dsic.upv.es

2 BACKGROUND

Earlier studies [13, 43] have revealed that a collective intelligence factor can emerge in human groups. We know that collectives can outperform individuals, and further that their performance is controlled by one or more of a) their organisational or network structure [29, 3, 30], b) the information aggregation details among their individuals [1], and c) the diversity between their members [20, 17]. Crowd-computing and crowd-sourcing [32, 24, 2] methodologies are excellent examples of CI that harness the wisdom of the crowd [37].

After carefully looking at the literature on collective intelligence including the abovementioned works and others including [28, 42, 38, 8, 41], we filter a set of factors or features from these works - that are not coupled to one particular cognitive system, problem or environment - which are intimately relevant to the performance of collectives, some of which are the number of members in a group, the communication or interaction protocol, as well as the difficulty of the environment. Curiously, there are some other factors which are often relatively neglected, such as the reasoning/learning speed of the agents and the interaction time of the collective as a whole. These features, in addition to some hypothetical combinations of them (grouped in ellipses) are depicted in Figure 1. It is not known *in which circumstances* and *how much* each these features individually influences the intelligence of the group, let alone the simultaneous influence of multiple features combined, which is what we attempt to *quantitatively* investigate in this paper.

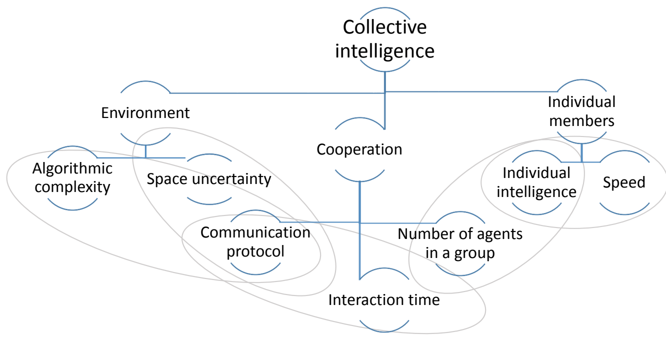

Figure 1: Factors and features relevant to the notion of collective intelligence (CI) perceived throughout various cognitive systems, and some hypothetical relationships between them grouped in ellipses.

The next section introduces our methodology for assessing individual and collective agent performances. The agent behaviours to be evaluated and their communication and interaction protocols are described in Sections 4 and 5 respectively. After we present our experimental setup in Section 6, we discuss and analyse our results from these experiments (in Sections 7 and 8) by making a series of observations on how the intelligence of the evaluated agents was influenced by a collection of factors, and draw some interesting conclusions connecting the research outcomes. We conclude in Section 9 by a brief summary and give some directions for future work.

3 EVALUATING INTELLIGENCE

To achieve our aims, we need a dynamic environment in which we can assess the influence of the factors appearing in Figure 1 on the performance of various types of cognitive systems over different environmental settings. While many environments could be appropriate, we have chosen for our purpose the *Anytime Universal Intelligence Test* (ANYNT) [18], which is derived from formal information theoretic backgrounds that have been practically used to evaluate *diverse kinds of entities* [21, 5, 6, 22], and was proven [18] to be an unbiased, dynamic setting which can be stopped at anytime.

3.1 The Anytime Universal Intelligence Test

We introduce the Λ^* (Lambda Star) environment class that focuses on a restricted - but important - set of tasks in AI. This environment extends the Λ environment class [18, Sec. 6][23] which implements the theory behind the Anytime Universal Intelligence Test [18]. The general idea is to evaluate an agent that can perform a set of actions, by placing it in a grid of cells with two special objects, *Good* ($\oplus$) and *Evil* ($\ominus$), travelling in the space using movement patterns of measurable complexities. Rewards are defined as a function of the position of the evaluated agent with respect to the positions of $\oplus$ and $\ominus$.

3.1.1 Structure of the test

We generate an environment space as an m-by-n grid-world populated with objects from $\Omega = \{\pi_1, \pi_2, \ldots, \pi_x, \oplus, \ominus\}$, the finite set of objects. The set of evaluated agents $\Pi \subseteq \Omega$ is $\{\pi_1, \pi_2, \ldots, \pi_x\}$. Each element in Ω can have actions from a finite set of actions $\mathcal{A}$ ={*left, right, up, down, up-left, up-right, down-left, down-right, stay*}. All objects can share the same cell at the same time except for $\oplus$ and $\ominus$ where in this case, one of them is randomly chosen to move to the intended cell while the other one keeps its old position. In the context of the agent-environment framework [25], a test episode consisting of a series of ϑ iterations is modelled as follows:

1. the environment space is first initialised to an m-by-n toroidal grid-world, and populated with a subset of evaluated agents from $\Pi \subseteq \Omega$, and the two special objects $\oplus$ and $\ominus$,
2. the environment sends to each agent a description of its range of 1 *Moore neighbour* cells [16, 40] and their contents, the rewards in these cells, as an observation,
3. the agents (communicate/interact and) respond to the observations by performing an action in $\mathcal{A}$, and the special objects perform the next action in their movement pattern,
4. the environment then returns a reward to each evaluated agent based on its position (distance) with respect to the locations of the special objects,
5. this process is repeated again from point #2 until a test episode is completed.

We are using a toroidal grid space in the sense that moving off one border makes an agent appear on the opposite one. Consequently, the distance between two agents is calculated using the surpassing rule (toroidal distance) such that, in a 5-by-5 grid space for example, the distance between cell $(1, 3)$ and $(5, 3)$ is equal to 1 cell.

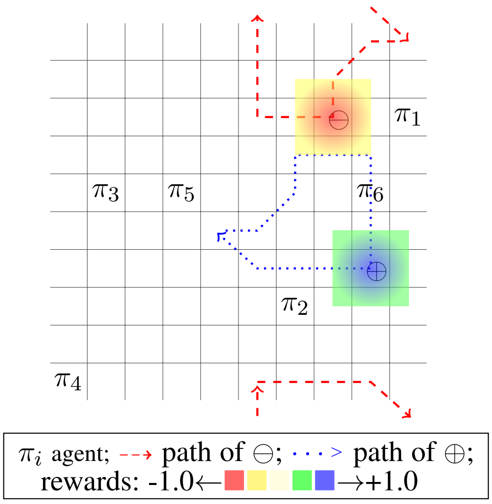

Figure 2: A diagrammatic representation of a sample 10-by-10 Λ^* environment, used to implement the theory behind the Anytime Universal Intelligence Test [18].

3.1.2 Rewarding function

The environment sends a reward to each evaluated agent from the set of rewards $\mathcal{R} \subseteq \mathbb{Q}$ where $-1.0 \leq \mathcal{R} \leq 1.0$.

Let $d(a, b)$ denote the (toroidal) distance between two objects a and b. Given an agent π_j, its positive reward at one test iteration is calculated as: $1/(d(\pi_j, \oplus) + 1)$ if $d(\pi_j, \oplus) < 2$, or 0 otherwise. Likewise its negative reward at that iterations is: $-1/(d(\pi_j, \ominus)+1)$ if $d(\pi_j, \ominus) < 2$, or 0 otherwise. Its total reward, r_j^i at iteration i, is the sum of its positive and negative rewards at that iteration.

3.2 Algorithmic Complexity

We regard the Kolmogorov complexity [27] of the movement patterns of the special objects as a measure of the algorithmic complexity $K(\mu)$ of the environment μ in which they operate. For instance, a Λ^* environment of high Kolmogorov complexity is sufficiently rich and structured to generate complicated (special object) patterns/sequences of seeming randomness. We measure the Lempel-Ziv complexity [26] of the movement patterns as an approximation to $K(\mu)$ as suggested in [26, 14]. Note that, at one test episode, the movement patterns of $\oplus$ and $\ominus$ are different but (algorithmically) equally complex. The recurrent segment of the movement pattern is at least of length one and at most $\lfloor \vartheta/2 \rfloor$, cyclically repeated until the final iteration (ϑ) of the test.

3.3 Search Space Complexity

We measure the search space complexity $\mathcal{H}(\mu)$ as the amount of *uncertainty* in μ, expressed by Shannon's entropy [35]. Let N be the set of all possible states of an environment μ such that a state s_μ, is the set holding the current positions of the special objects

$\{\oplus, \ominus\}$ in the m-by-n space. Thus the number of states $|N|$ increases with the increase in the space dimensions m and n, and it is equal to the number of permutation ${}^{\mathrm{m}\times\mathrm{n}}P_2 = \frac{(\mathrm{m}\times\mathrm{n})!}{(\mathrm{m}\times\mathrm{n}-2)!}$. The entropy is maximal at the beginning of the test as, from an agent's perspective, there is complete uncertainty about the current state of μ. Therefore $p(s_\mu)$ follows a uniform distribution and is equal to $1/|N|$. Using $\log_2$ as a base for our calculations, we end up with: $\mathcal{H}(\mu) = -\sum_{s_\mu \in N} p(s_\mu) \log_2 p(s_\mu) = \log_2 |N|$ bits.

Despite the test's being originally designed to return a general measure of intelligence, we do not make this assumption in this paper. Nevertheless, we appraise the test, at a minimum, as an accurate measure of the testee's ability of performing over a class of: inductive inference, compression and search problems, all of which are particularly related to intelligence [9, 10, 11, 19, 34, 12]. Note, however, that we will use the term *intelligence* to describe the effectiveness or accuracy of an evaluated agent over this test. It is of great importance that the illustrative class of problems assessed by the test is shared across, and applies to, various types of cognitive systems since this meets our criteria for the evaluation, as raised in the introduction.

4 AGENT TYPES AND BEHAVIOURS

We evaluated agents of five different behaviours, both in isolation and collectively in (cooperative) groups over the Λ^* environment. A description of these agents is given in the following paragraphs.

Local search agents: given an agent π_j, we denote by c_j^i and $r(c_j^i)$ the cell where π_j is located at iteration i, and the reward in this cell respectively. Let N_j^i and $R(N_j^i)$ denote respectively the set of range of 1 *Moore* neighbour cells [16, 40] of agent π_j (including c_j^i) at iteration i, and the reward values in these cells. $\mathtt{R}(c_j^i, a)$ is a function that returns the reward agent π_j gets after performing action $a \in \mathcal{A}$ when it is in cell c_j^i. The behaviour of local search agents is defined as follows: $a_j^i \leftarrow \arg\max_{a\in\mathcal{A}} \mathtt{R}(c_j^i, a)$. If all actions return an equivalent reward, then a random action in $\mathcal{A}$ is selected.

Reinforcement learning agents: two of the most frequently used RL (reinforcement learning) behaviours are Q-learning [39] and Sarsa [33, 39]. In the Q-learning behaviour, agents learn using an action-quality function in order to find the best action-selection policy for a given MDP (Markov Decision Process). Alternatively, Sarsa agents learn a MDP policy using an on-policy temporal-difference learning technique. Before learning starts, we initialise the elements of the Q-table to 2.0 so that the quality of a state-action pair, $Q \leftarrow S \times \mathcal{A}$, is always positive despite that rewards fall in the range $[-1.0, 1.0]$. Because the testing environment is dynamic, each state in S was designated to be the unique combination of one cell position c at one iteration i of the test, leading to a total number of states[3] $|S| = (\mathrm{m}\times\mathrm{n})^{\vartheta}$. Before evaluation, we trained the RL agents for 100 rounds previous to each episode using both a discount factor γ and a learning rate α of 0.30, selected after fine-tuning these parameters on a single agent scenario to reach a general (average) optimal payoff. Our agents learn offline, and thus cease to update their Q-table once their training is complete.

Oracle agents: an *oracle* agent knows the future movements of $\oplus$, the *Good* special object. At each step i of an episode this agent approaches the subsequent $i+1$ cell destination of $\oplus$ seeking maximum payoff. However, if $\oplus$ has a constant movement pattern (e.g., moves constantly to the right) pushing it away from the oracle, then the oracle will move in the opposite direction in order to intercept $\oplus$ in the upcoming test steps. Once it intercepts $\oplus$, it then continues operating using its normal behaviour.

Random agents: a random agent randomly choses an action from the finite set of actions $\mathcal{A}$ at each iteration until the end of an episode.

The scores of the random and oracle agents will be used as a baseline for our experiments, where a random agent is used as a lower bound on performance while the oracle is used as an upper bound.

5 COMMUNICATION PROTOCOLS

The agents were also evaluated collectively in groups. A description of the interaction and communication protocols used in these collectives are given below.

Stigmergy or indirect communication: we propose a simple algorithm for enabling communication between local search agents using stigmergy [15] (indirect communication). For instance, we let the agents induce fake rewards in the environment, thus indirectly inform neighbour agents about the proximity of the special objects. Note that fake rewards will not affect the score (real reward payoff) of the agents. Let $\hat{R}(N_j^i)$ denote the set of fake rewards in the neighbour cells of agent π_j (including c_j^i) at iteration i, and $\hat{\mathtt{R}}(c_j^i, a)$ is a function returning the fake reward agent π_j gets after performing action $a \in \mathcal{A}$ when it is in cell c_j^i at iteration i. Fake rewards are induced in the environment according to Algorithm 1. Each agent proceeds

Algorithm 1 Stigmergic or indirect communication: fake reward generation over one iteration i of the test.

Input: Π (set of evaluated agents), $0 < \gamma < 1$ (fake reward discounting factor), a test iteration i.
Initialize: $\forall \pi_j \in \Pi$: $\hat{R}(N_j^i) \leftarrow 0.0$.
Begin
  for j $\leftarrow 1$ **to** $|\Pi|$ **do** ▷ loop over agents
    $r^{max} \leftarrow \max R(N_j^i)$
    $r^{min} \leftarrow \min R(N_j^i)$
    $\hat{r} \leftarrow \gamma(r^{max} + r^{min})$ ▷ average expected reward
    $\hat{R}(N_j^i) \leftarrow R(N_j^i) + \hat{r}$
  end for
End

by selecting an action by relying on fake rewards this time instead of the real rewards, as follows: $a_j^i \leftarrow \arg\max_{a\in\mathcal{A}} \hat{\mathtt{R}}(c_j^i, a)$. If all actions are equally rewarding, then a random action is selected. Thereupon, we expect local search agents using stigmergy to form non-strategic coalitions after a few iterations of the test as a result of tracing the most elevated fake rewards in the environment.

Implicit leadership through auctions and bidding: in this cooperative setting, local search agents go into a *single dimensional English auction* [31] at each iteration i, and bid on the right to lead the other agents in their group by appointing one target cell to be approached. At each iteration, each auctioneer (agent) generates a value of the maximum reward existing in its neighbourhood, which is then used as its bidding "money" for the auction. The richest agent[4] wins the auction visibly to all the other agents. It then selects the *target cell* to be approached by all other agents in the collective. This bidding behaviour is described in Algorithm 2 in which $n_j^i \in N_j^i$ and $r(n_j^i)$ denote one of the *Moore* neighbour cells of agent π_j (without excluding c_j^i) at iteration i, and the reward in this cell respectively.

Imitating super-solver agents: a group of isolated local search agents is put in the same space with one (unevaluated) oracle agent.

[3] Recall that m and n refer to the grid space dimensions, while ϑ is the number of iterations in a single test episode.

[4] If more than one agent are equally rich then, for the sake of simplicity, the last one to participate in the auction wins.

Algorithm 2 Single dimensional English auction at one iteration i of the test.

Input: Π (set of evaluated agents), -1.0 <bid< 1.0, a test iteration i.
Initialize: bid $\leftarrow -1.0$
Begin
  for j $\leftarrow$ 1 **to** $|\Pi|$ **do** $\triangleright$ loop over agents
    $money \leftarrow \max R(N_j^i)$
    if $money >= bid$ **then**
      $bid \leftarrow money$
      $target \leftarrow \arg\max_{n_j^i \in N_j^i} r(n_j^i)$ $\triangleright$ set the target to the neighbour cell n_j^i holding the highest reward $r(n_j^i)$ at iteration i
    end if
  end for
End

Local search agents imitate the oracle by following it into the same cell only when it is in their visibility range (neighbourhood) otherwise, they operate using their normal behaviour.

Wisdom of the crowd (WOC) by information aggregation: where the collective opinion of the evaluated agents is aggregated from the opinions of all the members of the collective.

In the case of reinforcement learning collectives, we let their members share and update a common Q-table, thus making them all learn and coordinate simultaneously. We evaluated both Q-learning and Sarsa collectives independently.

In the case of local search collectives, the observations of all agents in the collective are aggregated into one global observation (and rewards from these observations are averaged in the case of overlap). Then, each member proceeds by selecting the action maximising its reward in line with the global observation.

6 EXPERIMENTAL SETUP

Each experiment consists of 1000 episodes (runs) of the test, each consisting of a number of iterations equal to 50. In each episode, agents are administered over a different task with complexity $K(\mu)$, such that $K(\mu) \in [2, 23]$, where a $K(\mu)$ of 23 corresponds to, a more or less, complex pattern prediction or recognition task. Moreover, in each episode, the collectives are re-initialised with different spatial (network) arrangements between their members.

Local search agents were evaluated in isolation as well as collectively using four communication or interaction protocols: stigmergy, implicit leadership, imitation (of the oracle agent) and harnessing the WOC through information aggregation. Likewise, reinforcement learning agents were evaluated in isolation and collectively by harnessing the wisdom of the crowd (WOC) through sharing and updating a common Q-table.

Test experiments were conducted over different search space uncertainties $\mathcal{H}(\mu)$, and the (intelligence) scores (in the range $[-1.0, 1.0]$) of the evaluated agents/collectives averaged over the 1000 episodes were recorded. The score of the collective is calculated as the mean of the scores of its members. For instance, the metric of (individual agent) universal intelligence defined in [18, Definition 10] was extended into a collective intelligence metric (Definition 2) returning an average reward accumulation per-agent measure of success (Definition 1) for a group of agents Π, over a selection of Λ^* environments.

Definition 1 *Given a Λ^* environment μ and a set of (isolated or interactive) agents $\Pi = \{\pi_1, \pi_2, \ldots, \pi_n\}$ to be evaluated, the (average per-agent per-iteration) reward $\tilde{R}_{\Pi,\mu,\vartheta}$ of Π over one test episode of ϑ iterations is calculated as:* $\tilde{R}_{\Pi,\mu,\vartheta} = \frac{\sum_{j=1}^{n}\sum_{i=1}^{\vartheta} r_j^i}{n\vartheta}$.

Definition 2 *The (collective) intelligence of a set of agents Π is calculated as: $\frac{1}{\omega}\sum_{\mu \in L} \tilde{R}_{\Pi,\mu,\vartheta}$, where L is a set of ω environments $\{\mu_1, \mu_2, \ldots, \mu_\omega\}$ such that $\forall \mu_t, \mu_q \in L$: $\mathcal{H}(\mu_t) = \mathcal{H}(\mu_q)$, and $\forall \mu_i \in L, K(\mu_i)$ is extracted from a range of (special object movement patterns with) algorithmic complexities in $]1, K_{max}]$.*

7 RESULTS AND DISCUSSION

Table 1: Intelligence test scores for collectives of 10 agents across different environment uncertainties $\mathcal{H}(\mu) \in [13.2, 19.6]$ bits, evaluated for 50 test-iterations. A plot of these results is also found in Figure 3.

	$\mathcal{H}(\mu)$ **value in bits**	**13.2**	**15.6**	**17.2**	**18.5**	**19.6**
1	Random agent	*-0.00079*	*0.00048*	*0.00008*	*-0.00013*	*0.00002*
2	Local search (LS) agent	*0.3365*	*0.1696*	*0.0936*	*0.0575*	*0.0423*
3	LS collective using stigmergy	*0.4025*	*0.2555*	*0.1431*	*0.0829*	*0.0579*
4	LS collective harnessing the WOC	*0.3828*	*0.3475*	*0.3118*	*0.2601*	*0.2110*
5	LS collective using implicit leadership	*0.3744*	*0.2842*	*0.2143*	*0.1722*	*0.1438*
6	LS collective using imitation	*0.5729*	*0.2880*	*0.1666*	*0.1022*	*0.0731*
7	Q-learning agent	*0.2516*	*0.0950*	*0.0484*	*0.0301*	*0.0207*
8	Q-learning collective harnessing the WOC	*0.4030*	*0.1832*	*0.0870*	*0.0482*	*0.0309*
9	Sarsa agent	*0.2708*	*0.1007*	*0.0501*	*0.0308*	*0.0228*
10	Sarsa collective harnessing the WOC	*0.4511*	*0.2042*	*0.1010*	*0.0563*	*0.0348*
11	Oracle agent	*0.8207*	*0.7905*	*0.7619*	*0.7339*	*0.7059*

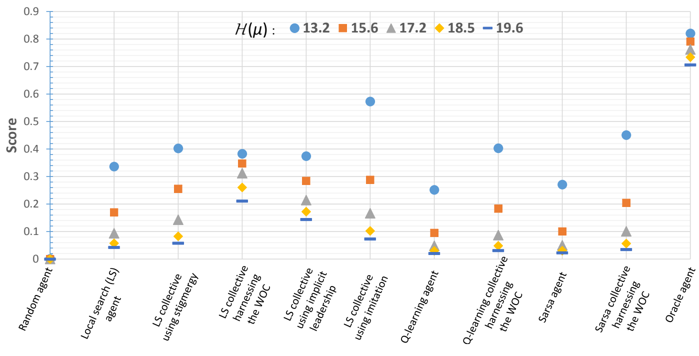

Figure 3: A plot of the test scores appearing in Table 1.

Sample results from the previously-mentioned experiments run in different environment (search space) uncertainties $\mathcal{H}(\mu)$ are listed in Table 1 for isolated agents and collectives Π, each having a number of agents or members $|\Pi| = 10$ agents. The standard deviation of the test scores σ is less than 0.001 between identical experiments.

7.1 Collectives outperform individuals

Results in Figure 3 clearly show (in at least three separate cases) that cooperative or interactive individuals can be more effective than isolated ones. (From Definition 1, the score of the whole is more than the sum of its parts.) This is consistent with earlier results (e.g., [29, 1]) for obvious reasons owing to diffusion of information (synergy) leading to the reduction of uncertainty inside the collective.

Yet the question remains, what are the dynamics which have led to such results? We recall our main aims, which consist of investigating and quantifying the influence of a list of factors on (individual and collective) intelligence. We address each of these factors in detail in the remainder of this paper.

7.2 Communication and interaction protocol

We observe in Figure 3 that the effectiveness of the (same selection of) agents is highly dependent on the collective decision-making technique or the communication protocol used to aggregate the information received from these agents. For instance, adopting auctions in local search collectives to claim leadership can be more effective than using stigmergy over some settings. Figure 3 also shows that, under certain circumstances, introducing heterogeneity in a group of local search agents by imitating a (super-solver) oracle agent leads to more effective coalitions that outperform their homogeneous (and isolated) peers by aggregating new information into the collective. However, the comparison between local search collectives is rather more complicated as their intelligence measures seem to further depend on the uncertainty of the testing environment, and not only on the interaction protocol. We also observe that harnessing the wisdom of the crowd by aggregating the observations of local search agents is very effective over highly uncertain environments, yet not exceptionally efficient in the opposite situation. The latter protocol seems to be very robust (in comparison to others) with respect to the changes in the uncertainty of the search space. A more thorough analysis on the efficiency of the examined communication protocols over different problem uncertainties is addressed in the following subsection.

In the case of RL agents, we observe that agents of different types (Q-learning and Sarsa) using the same cooperation technique to aggregate their information have achieved different scores. Sarsa agents outperform Q-learning agents up to about a similar extent both in cooperative and isolated settings. This indicates that the collective intelligence of the group also depends on, and is correlated with, the individual intelligence (or the type) of the agents in the group.

Furthermore, despite the broad differences in the interaction protocols and the wide range of task complexities, collective intelligence manifested across the various collectives, showing that CI can also emerge in a non-human context or environment, thus reinforcing and adding to the conclusions of [13] conclusions.

7.3 Uncertainty in the environment

Figure 3 shows that the performance of the evaluated agents decreases with the increase in uncertainty[5] $\mathcal{H}(\mu)$, in accordance with former tests [22] that have been applied on humans and artificial agents. Moreover, the gap between the scores of the isolated and cooperative agents varies in view of the uncertainty in the environment, but the relationship between both variables cannot be easily grasped from the figure.

We wish to measure the variation in the weight of the cooperative agents' scores to their score in the isolated setting, across different environment uncertainties. Therefore, we define the *coefficient of effectiveness* $\theta = \alpha/\beta$, as the ratio of the score of a set of agents Π working in some cooperative scenario ($\alpha = score(\Pi^{coop})$), to its score in the isolated scenario $\left(\beta = score(\Pi^{isolated})\right)$. We calculated θ for the different agent types across different uncertainties and plotted the results in Figure 4.

Figure 4a shows that the θ values corresponding to local search groups using imitation ($\theta^{imitation}$) and those relying on stigmergy ($\theta^{stigmergy}$) are more or less steady across the different $\mathcal{H}(\mu)$ values, implying that each of these protocols is approximately equally advantageous over different problem uncertainties. In addition, $\theta^{imitation} > \theta^{stigmergy}$ over the selected uncertainties, and thus collectives relying on imitation are more effective than those using stigmergy. The observations are more interesting for collectives using auctions to claim leadership. For instance, in environments of uncertainties lower than 16 bits, imitating a smart agent is more advantageous than following a leader. Whereas, the inverse is true for environments of higher uncertainties. The effectiveness of local search agents using auctions significantly increases to become much higher than that of the same group of agents imitating an oracle. Similar results are observed for local search collectives harnessing the wisdom of the crowd. We conclude that relying on the best (super-solver) agent in the groups does not guarantee an optimal performance. This is somewhat consistent with [20]'s claims re diversity vs. ability, even though the intuitions here are different. These results have a fundamental impact on the choice of the communication protocol to be used in order to aggregate the information received from a group of agents, especially over problems where the search complexity can be estimated in advance.

For RL agents, Figure 4b shows an overall similar shift in effectiveness for both Q-learning and Sarsa collectives. Their performances significantly increase over isolated agents to reach a peak around $\mathcal{H}(\mu) = 16$ bits, but then start to drop down over higher uncertainties. This illustrates the fact that cooperative RL collectives are most advantageous over environments which are somewhat highly-uncertain for isolated RL agents to be efficient, yet not too uncertain for them (the cooperative collectives) to be considerably efficient. In other words, collective intelligence might only be slightly perceived in groups operating in very simple environments (or problems) where individuals could perform relatively well, or in those too difficult (broad) to be explored within a limited *interaction time*, given a limited *number of members* in the group. Thus, in order to understand the global picture of collective behaviour and its dynamics, it is crucial to look into the latter two factors (interaction time and number of members) and measure their effects, if any exist, on group intelligence.

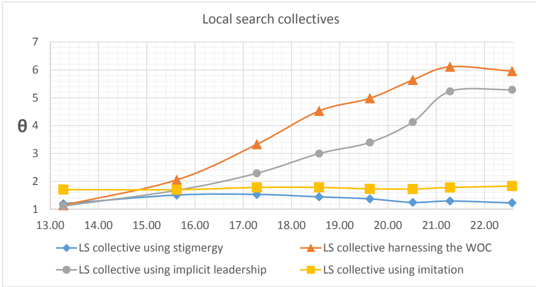

(a) Local search (LS) collectives.

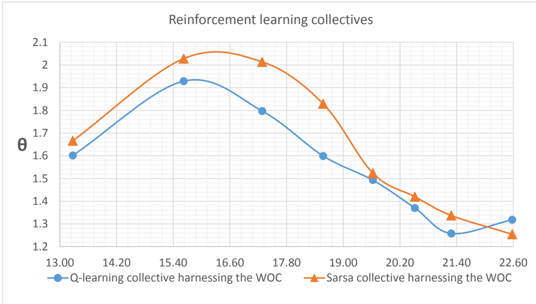

(b) Reinforcement learning (RL) collectives.

Figure 4: Shift in effectiveness θ for local search and RL agents over different environment uncertainties in bits.

[5] Except for random agents which always score around zero.

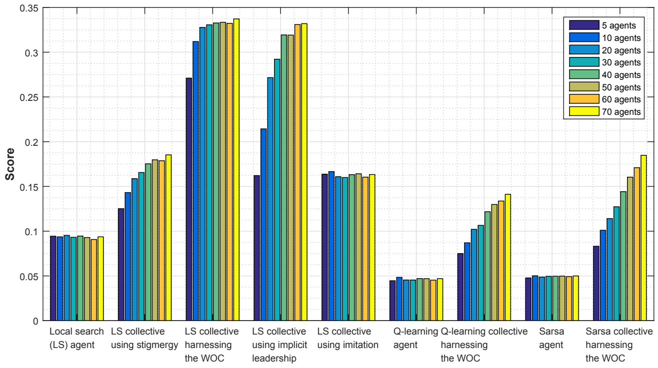

Figure 5: Intelligence scores recorded across different numbers of agents $5 \leq |\Pi| \leq 70$, in 17.8-bit $\mathcal{H}(\mu)$ environments.

7.4 Number of agents in a group

In all our previous experiments the number of evaluated agents in each collective was set to 10. Whereas, Figure 5 illustrates the scores of the evaluated collectives across different number of agents varying between 5 and 70. The general picture shows that local search collectives relying on stigmergy and auctions gradually improve in performance as more agents are added into the collective. This is not the case for collectives relying on imitation, which only show a shallow variation in score. In fact, local search agents relying on imitation performed better than those using auctions when $|\Pi|$ was set to 5 agents. However, the opposite was true when we increased the number of agents to 10 and higher. This illustrates that, when the group is small in number, relying on a super-solver agent might be more advantageous than interacting between the individual members, however, as the group gets larger, more information is added into the collective and the expertise of a single oracle becomes rudimentary in comparison to the aggregated experiences (synergy) from individual members. Moreover, we observe that local search collectives harnessing the WOC improve faster in performance than those following a leadership. Nonetheless, when the number of agents gets higher the performances of these two collectives get closer to one another.

We also observe that, increasing the number of local search agents is more effective and has greater influence on the scores for agents relying on auctions than those using stigmergy. On another hand, the increase in efficiency is slightly non-linear to the number of agents introduced. For instance, the main improvements in scores are more concentrated at the early introductions of agents. Afterwards the scores continue to rise, but less and less significantly.

Similar observations illustrate that the effectiveness of RL collectives harnessing the WOC improves as we increment the number of agents. Moreover, Sarsa collectives seem to be slightly more efficient than Q-learning collectives as new agents are introduced into the group. The key issue in this experiment is that, collective intelligence cannot be considered independently of the number of members in the group. Instead, it is a function of - so far - at least three factors, each having a different influence that we have measured, and bearing distinctive properties of which we have identified some.

7.5 Time and intelligence

In this paragraph we address the relevance of time to intelligence, which is often ignored in the assessment of collective intelligence. Figure 6 shows the variations in the intelligence scores as we extend

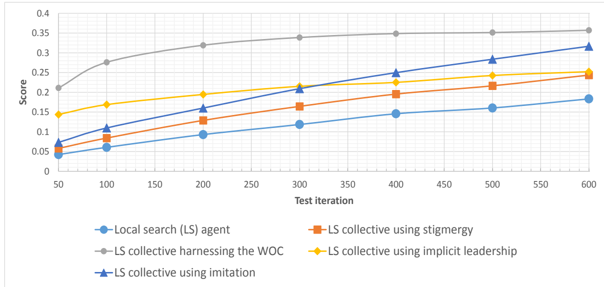

(a) Local search agent collectives in $\mathcal{H}(\mu) = 17.2$ bits.

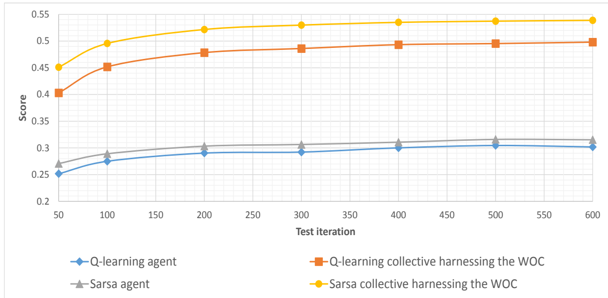

(b) RL collectives in $\mathcal{H}(\mu) = 13.2$ bits.

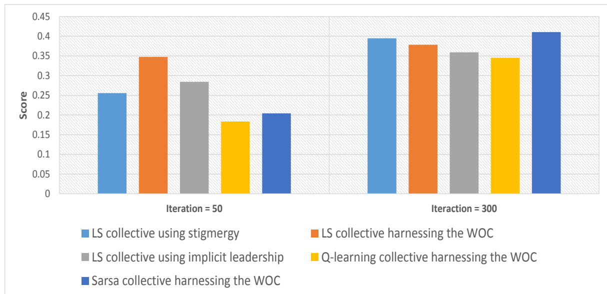

(c) Scores in $\mathcal{H}(\mu) = 13.2$ bits after 50 and 300 iterations.

Figure 6: Variations in the agents' intelligence scores as we extend the evaluation time (number of iterations) of the test.

the interaction time (number of iterations or interaction steps) of the test. We observe that some scores incline to converge as more time is given to the members to perform on the test. Figure 6a shows that the advantage of cooperative local search agent groups over isolated agents is higher at the early stages of the test in the case of agents using auctions to claim leadership. Afterwards, the gap in performance slowly decreases with time until iteration number 600. On the contrary, the gap between the scores of local search agents imitating an oracle and their isolated peers grows as we let the test run, implying that local search agents relying on imitation require longer periods of time to reach their best performance. We have already shown in Figure 3 that over some uncertainties, local search agents relying on auctions outperform those imitating an oracle, which is again consistent with the results in Figure 6a (up to iteration 300). However, this experiment also suggests that imitating a super-solver is highly rewarding over time, leading to better-scoring collectives than when using leadership through auctions. These results illustrate how diverse social organisations between the members of the collective determine its performance over time. For instance, a (dynamic) leadership scheme or organisation seems to be more rewarding than a simple flat hierarchy relying on stigmergic communication given a limited interaction time with the environment.

Moreover, the general picture shows that a local search collective harnessing the WOC is most advantageous over isolated agents (and other collectives) mainly before the 300^{th} iterations, at which point its performance begins to converge slowly.

In the case of RL agents (Figure 6b), both isolated agents and collectives improve in performance with time, keeping an overall steady relationship between the differences in their scores. This raises another concern re the intelligence of artificial agents. It is intriguing as to what ideally counts as more intelligent, a fast re-active agent with a humble performance, or a slow one with an exceptional performance over an extended period of time? Should we consider the *potential intelligence* of an agent instead? To understand the importance of time in measuring intelligence, we compare the scores of RL and local search collectives over 13.2-bit $\mathcal{H}(\mu)$ environments after 50 and 300 test iterations as illustrated in Figure 6c. We find that local search collectives outscored Sarsa collectives up to the first 50 iterations while the opposite is true at iteration 300. This type of experiment is one of the most revealing of how the (communication and interaction) reasoning/learning speed of multiagent systems influences their measured performance given a finite/bounded operation or interaction time.

7.6 Algorithmic complexity and intelligence

In this paragraph we shed some light on how the performance of (groups of) agents is influenced by the algorithmic complexity of the task. To minimise the effect of search and exploration (relative to exploitation) on the scores we initialised all agents to neighbour locations from $\oplus$. We then evaluated the agents over tasks of different algorithmic complexities (randomness) $K(\mu)$ grouped into three difficulty levels: easy $\in [6, 8]$, medium $\in [9, 13]$ and hard $\in [14, 19]$. This experiment stands out from previous related experiments in the field, as collectives are assessed against tasks of quantifiable algorithmic complexity, as opposed to ones qualitatively ranked based of their difficulty. Results illustrated in Figure 7 show that the performance of artificial agents, similar to that of individual human performance [22], decreases when evaluated over patterns of higher algorithmic complexities. For instance, learning and predicting random patterns is more difficult, per se, than learning or inferring compressible ones. Moreover, this experiment suggests that RL collectives are better learners than their isolated peers since the difference between the cooperative agents' scores over the 3 levels of difficulties is significantly smaller than that of the isolated ones. What's more intriguing in Figure 7 is the difference in behaviour between cooperative RL agents and local search collectives. While RL collectives are still more effective over isolated agents when evaluated over learning problems, all local search agents (isolated and collectives) performed equally when the effect of *search and exploration* was minimised. More importantly, we find that RL collectives are more robust with respect to the change in algorithmic complexity as opposed to local search agents which display a wide gap in scores over the three levels of complexity.

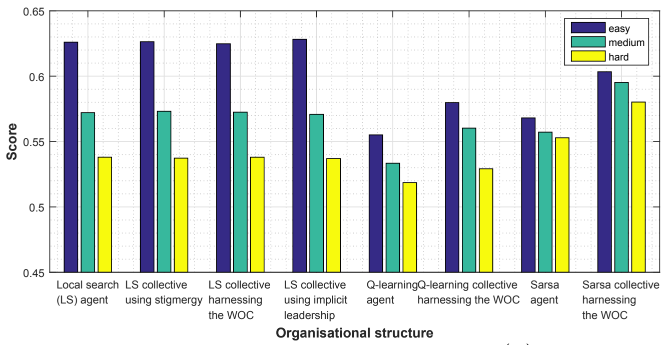

Figure 7: Scores over different task complexities $K(\mu)$ using collectives of $|\Pi| = 10$ agents, evaluated in 13.2-bit $\mathcal{H}(\mu)$ environments for 50 interactions.

All in all, what this experiment suggests is that, further to the previously examined factors, (collective) intelligence is a function of the agent type and the algorithmic complexity of the given task, both combined.

8 ORGANISATIONAL BEHAVIOUR

In spite of the different communication protocols we have evaluated, it is still not clear how the organisational structure of the group [3, 30] affects its performance on intelligence tests. Therefore, we have further evaluated the performance of equally sized collectives of local search agents organised in four different (divisional and network) structures and studied their organisational behaviour. These structures are illustrated in Figure 8 below. In the flat, fully connected, structure (Figure 8a) all agents share their observations between one another. This absolute aggregation of information leads to a similar effect as that of local search collectives harnessing the wisdom of the crowd. In the subgroup structure (Figure 8b) we divide the collective into four smaller subgroups. Each one of those subgroups then implements a flat structure as the one described previously. In the hierarchical structure (Figure 8c), each (non-leaf) agent receives feedback from its children at each iteration of the test before selecting an action. Leaf-nodes operate in isolation. Finally, in the autocratic structure (Figure 8d), a single agent controls the actions of the rest of the collectives irrespective of its members' observations.

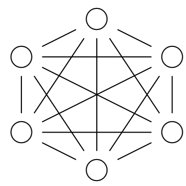

(a) Flat (fully-connected) structure.

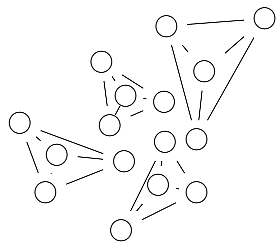

(b) Subgroup structure.

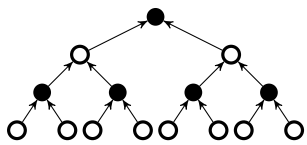

(c) Hierarchical (tree-like) structure.

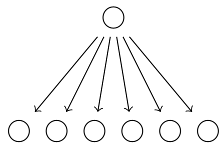

(d) Autocracy.

Figure 8: Graphical representation of different group organisational structures. Nodes represent agents and edges reflect the flow of communication and interaction between these agents.

The results from our experiments show that flat, fully-connected, network structures are the most efficient since they maximise the aggregation of information received from the members of the collec-

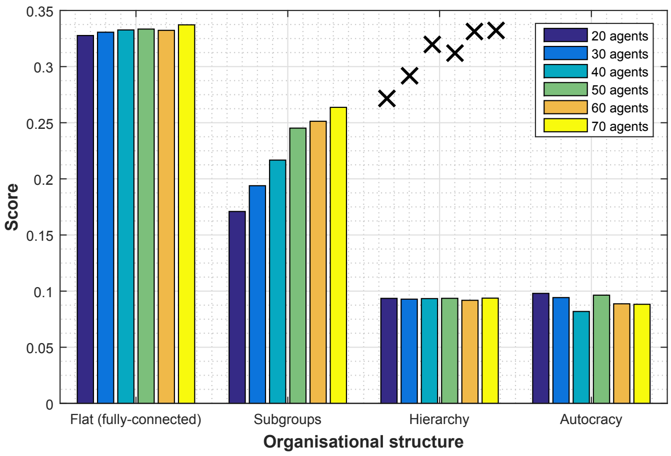

Figure 9: Scores of local search collectives organised in different network structures, across various number of agents. The collectives are evaluated in 17.2-bit $\mathcal{H}(\mu)$ environments for 50 iterations. The label $(\times)$ in this figure depicts the average score of the agent at the top of the hierarchy (root node of the binary-tree).

tive. However, it is known that this type of structure is very costly as it requires a large number of connections ($n(n-1)/2$ connections, where n is the number of agents) to be introduced between the members of the group [36]. In this type of organisation (and for the corresponding environment uncertainty and evaluation-time parameters) the number of agents did not significantly affect the performance of the collective. Whereas, after dividing this collective into smaller subgroups, the number of agents turns out to be of major importance[6]. The effectiveness of each subgroup improved gradually with the increase in the number of agents thus reducing the gap in performance between this organisational structure and fully-connected one. This shows that dividing a collective into smaller groups is most beneficial for highly populated collectives, especially when the number of connections inside the collective grows very large and becomes a bottleneck on communication.

In the hierarchical and autocratical structures, the measured effectiveness is low compared to the previous two models. We observe that the average performance of a hierarchical group is slightly steadier than that of a group governed by single agent with absolute control on decision-making. Interestingly, we have noticed that in the hierarchical structure, high-scoring agents are the ones at the top of the hierarchy since (in our model) they receive feedback from their children (which in turn receive feedback from theirs) while the ones at the bottom perform in isolation and have low scores. Figure 9 shows that the average scores of the root agents in the hierarchy are significantly higher than the average score of the collective, indicating a high standard deviation between the members' scores in this organisational structure. Since the number of leaves is almost half the number of nodes[7], and the number of agents declines quickly as we move up the hierarchy, this organisation does not deliver a high average group performance.

Finally, our results show that the performance of a local search collective implementing an autocracy is similar on average to that of isolated local search agents. Agents in this organisational structure do not show any significant discrepancies in their scores or behaviours.

[6] Note that all four evaluated subgroups showed a similar performance, but scores were only plotted for the first subgroup to enhance readability.

[7] Number of leaves in a full binary tree is equal to $(\#nodes + 1)/2$.

9 CONCLUSIONS AND FUTURE WORK

We have addressed the relevance of several factors and their interaction to the notion of intelligence and its emergence. We first started by looking at the different contexts in which collective intelligence has been shown to emerge, from face-to-face human groups, group collaborations via the web, social insect colonies and swarms, etc. Accordingly, we filtered a series of factors and features that are not coupled to one particular cognitive system, problem or environment, and illustrated how they influence the collective behaviour of the group, and hinder its intelligence.

The studied factors were shown to have a major influence on the performance of collectives that we have also measured. But, what made our conclusions more intriguing is the peculiar nature of collective intelligence seen as a function of all the examined factors simultaneously, as well as some of them combined. We identified circumstances where one cooperative system outperformed another under some values or setups of the studied factors yet failed to do so under others (e.g., in Section 7.5, limited vs. extended interaction time and, in Section 7.3, low vs. high environment uncertainty), reflecting on how these factors independently but also jointly shape the effectiveness of multiagent systems, and the spread of intelligence in these systems. Some of our conclusions (in Section 7.3) reflected how relying on an expert (super-solver) agent in the group does not guarantee its optimal performance. We also measured the effect of introducing more agents into the group (Section 7.4), and showed that it is tightly controlled by the communication protocol used between its members. We have highlighted scenarios (in Section 7.6) where only some types of collectives outperform their equally sized group of isolated agents over (algorithmically) complex environments, and shown how the influence of the environment difficulty (uncertainty and complexity) is a major factor controlling the capacity for intelligence. Moreover, we looked (in Section 8) into how the effectiveness of (the same selection of) agents adopting different organisational and network structures can significantly vary from one structure to another.

We have answered a fundamental question by showing the existence, and *quantitatively* measuring the influence, of some general factors and principles shaping the spread of intelligence that are regularly perceived across different cognitive systems.

In order for our results to be transferred to a guideline for designing multiagent cooperation, we have released the source code and scripts to run our experiments as open-source in [7, Section 5.1]. This will allow both additional testing and extensions to (the current version of) the Λ^* environment. The motivation is to encourage people in the AI community to quantitatively evaluate new types of heuristics, algorithms, communication protocols and network structures.

Another future goal is to further evaluate agents and collectives over a wide range of general AI problems. For example, agents (isolated or collectives) could be evaluated over exploration/exploitation problems in an environment consisting of a hidden fitness landscape with many local, and only one global, optima as further elaborated in [7, Section 6.2]. Other possible examples might include pattern recognition (and sequence completion) problems, in which payoff is determined by how accurately a subject learns and predicts a pattern. Other general multiagent problems that require coordination [7, Section 6.1] (e.g., lifting and moving an object), or scheduling [4] (e.g., job shop scheduling), can be used for alternative evaluation.

Finally, we hope that - by following this direction in the quantification of intelligence - we would pave the way towards a rigorous and unified model of collectively intelligent groups and societies.

REFERENCES

[1] Luís M. A. Bettencourt, 'The rules of information aggregation and emergence of collective intelligent behavior', *Topics in Cognitive Science*, **1**(4), 598–620, (2009).

[2] Eric Bonabeau, 'Decisions 2.0: The Power of Collective Intelligence', *MIT Sloan Management Review*, (Winter), 45–52, (2009).

[3] John Child, 'Organizational structure, environment and performance: The role of strategic choice', *sociology*, **6**(1), 1–22, (1972).

[4] Nader Chmait and Khalil Challita, 'Using simulated annealing and ant-colony optimization algorithms to solve the scheduling problem', *Computer Science and Information Technology*, **1**(3), 208–224, (2013).

[5] Nader Chmait, David L. Dowe, David G. Green, and Yuan-Fang Li, 'Observation, communication and intelligence in agent-based systems', in *Proc. 8th Int. Conf. Artificial General Intelligence, Berlin, Germany*, volume 9205 of *Lecture Notes in Artificial Intelligence (LNAI)*, pp. 50–59. Springer, (Jul 2015).

[6] Nader Chmait, David L. Dowe, David G. Green, Yuan-Fang Li, and Javier Insa-Cabrera, 'Measuring universal intelligence in agent-based systems using the anytime intelligence test', Technical Report 2015/279, FIT, Clayton, Monash University, (2015).

[7] Nader Chmait, Yuan-Fang Li, David L. Dowe, and David G. Green, 'A dynamic intelligence test framework for evaluating AI agents', in *Proc. Evaluating General-Purpose AI (EGPAI), ECAI workshop*, (2016).

[8] Iain D. Couzin, 'Collective cognition in animal groups', *Trends in Cognitive Sciences*, **13**(1), 36–43, (2009).

[9] David L. Dowe and Alan R. Hajek, 'A computational extension to the Turing Test', *Proc. 4th Conf. of the Australasian Cognitive Science Society, University of Newcastle, NSW, Australia*, (1997).

[10] David L. Dowe and Alan R. Hajek, 'A computational extension to the Turing Test', *Technical Report #97/322, Dept Computer Science, Monash University, Melbourne, Australia*, (1997).

[11] David L. Dowe and Alan R. Hajek, 'A non-behavioural, computational extension to the Turing Test', in *International conference on computational intelligence & multimedia applications (ICCIMA'98), Gippsland, Australia*, pp. 101–106, (1998).

[12] David L. Dowe, José Hernández-Orallo, and Paramjit K. Das, 'Compression and intelligence: Social environments and communication', in *Proc. 4th Int. Conf. on AGI*, pp. 204–211, Berlin, (2011). Springer.

[13] David Engel, Anita W. Woolley, Ishani Aggarwal, Christopher F. Chabris, Masamichi Takahashi, Keiichi Nemoto, Carolin Kaiser, Young Ji Kim, and Thomas W. Malone, 'Collective intelligence in computer-mediated collaboration emerges in different contexts and cultures', in *Proc. 33rd Conf. on CHI*, pp. 3769–3778, NY, USA, (2015). ACM.

[14] Scott C. Evans, John E. Hershey, and Gary Saulnier, 'Kolmogorov complexity estimation and analysis', in *6th World Conf. on Systemics, Cybernetics and Informatics*, (2002).

[15] Plerre-P. Grassé, 'La reconstruction du nid et les coordinations interindividuelles chez Bellicositermes natalensis et Cubitermes sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs', *Insectes sociaux*, **6**(1), 41–80, (1959).

[16] Lawrence Gray, 'A mathematician looks at Wolfram's new kind of science', *Notices of the American Mathematical Society*, **50**(2), 200–211, (2003).

[17] Vahid Hashemi and Ulle Endriss, 'Measuring diversity of preferences in a group.', in *ECAI*, pp. 423–428, (2014).

[18] José Hernández-Orallo and David L. Dowe, 'Measuring universal intelligence: Towards an anytime intelligence test', *Artificial Intelligence*, **174**(18), 1508–1539, (December 2010).

[19] José Hernández-Orallo and Neus Minaya-Collado, 'A formal definition of intelligence based on an intensional variant of Kolmogorov complexity', in *Proc. of the Int. Symposium of EIS*, pp. 146–163. ICSC Press, (1998).

[20] Lu Hong and Scott E. Page, 'Groups of diverse problem solvers can outperform groups of high-ability problem solvers', *PNAS*, **101**(46), 16385–16389, (2004).

[21] Javier Insa-Cabrera, José-Luis Benacloch-Ayuso, and José Hernández-Orallo, 'On measuring social intelligence: Experiments on competition and cooperation', in *Proc. 5th Conf. on AGI*, eds., Joscha Bach, Ben Goertzel, and Matthew Iklé, volume 7716 of *LNCS*, pp. 126–135. Springer, (2012).

[22] Javier Insa-Cabrera, David L. Dowe, Sergio España-Cubillo, M. Victoria Hernandez-Lloreda, and José Hernández-Orallo, 'Comparing humans and AI agents.', in *AGI*, volume 6830 of *LNCS*, pp. 122–132. Springer, (2011).

[23] Javier Insa-Cabrera, José Hernández-Orallo, David L. Dowe, Sergio España, and M. Victoria Hernández-Lloreda, 'The ANYNT project intelligence test: λ_{one}', in *AISB/IACAP 2012 Symposium "Revisiting Turing and his Test"*, pp. 20–27, (2012).

[24] Mark Klein and Ana Cristina B. Garcia, 'Crowd computing: From human computation to collective intelligence (tutorial)', in *Proc. 24th IJCAI, Buenos Aires, Argentina, July*, (2015).

[25] Shane Legg and Marcus Hutter, 'Universal intelligence: A definition of machine intelligence', *Minds and Machines*, **17**(4), 391–444, (2007).

[26] Abraham Lempel and Jacob Ziv, 'On the Complexity of Finite Sequences', *Information Theory, IEEE Transactions on*, **22**(1), 75–81, (January 1976).

[27] Ming Li and Paul Vitányi, *An introduction to Kolmogorov complexity and its applications (3rd ed.)*, Springer-Verlag New York, Inc., 2008.

[28] Thomas W. Malone and Michael S. Bernstein, *Handbook of collective intelligence*, MIT Press, 2015.

[29] Winter Mason and Duncan J. Watts, 'Collaborative learning in networks', *PNAS*, **109**(3), 764–769, (2012).

[30] Henry Mintzberg, 'The structuring of organizations: A synthesis of the research', *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*, (1979).

[31] Simon Parsons, Juan A. Rodriguez-Aguilar, and Mark Klein, 'Auctions and bidding: A guide for computer scientists', *ACM*, **43**(2), 1–59, (Feb 2011).

[32] Massimo Poesio, Jon Chamberlain, Udo Kruschwitz, Livio Robaldo, and Luca Ducceschi, 'Phrase detectives: Utilizing collective intelligence for internet-scale language resource creation (extended abstract)', in *Proc. 24th IJCAI, Buenos Aires, Argentina*, pp. 4202–4206, (2015).

[33] G. A. Rummery and M. Niranjan, 'On-line Q-learning using connectionist systems', Technical Report TR 166, Cambridge University Engineering Department, Cambridge, England, (1994).

[34] Pritika Sanghi and David L. Dowe, 'A computer program capable of passing I.Q. tests', in *Proc. of the Joint International Conference on Cognitive Science, 4th ICCS International Conference on Cognitive Science & 7th ASCS Australasian Society for Cognitive Science (ICCS/ASCS-2003)*, ed., P. P. Slezak, pp. 570–575, Sydney, Australia, (13-17 July 2003).

[35] Claude E. Shannon, 'A mathematical theory of communication', *Bell System Technical Journal, The*, **27**(3), 379–423, (July 1948).

[36] Steven H Strogatz, 'Exploring complex networks', *Nature*, **410**(6825), 268–276, (2001).

[37] James Surowiecki, *The Wisdom of Crowds*, Anchor, 2005.

[38] Kagan Tumer and David Wolpert, 'A survey of collectives', in *In collectives and the design of complex systems*, pp. 1–42. Springer, (2004).

[39] Christopher J. C. H. Watkins and Peter Dayan, 'Technical note: Q-learning', *Mach. Learn.*, **8**(3-4), 279–292, (May 1992).

[40] Eric W. Weisstein. Moore neighborhood, from mathworld - a wolfram web resource. `http://mathworld.wolfram.com/MooreNeighborhood.html`, 2015. Last accessed: 2015-11-10.

[41] Michael Wooldridge, *An Introduction to MultiAgent Systems*, Wiley Publishing, 2nd edn., 2009.

[42] Anita W. Woolley, Ishani Aggarwal, and Thomas W. Malone, 'Collective intelligence and group performance', *Current Directions in Psychological Science*, **24**(6), 420–424, (2015).

[43] Anita W. Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone, 'Evidence for a collective intelligence factor in the performance of human groups', *Science*, **330**(6004), 686–688, (2010).