Efficient Opinion Sharing in Large Decentralised Teams

Oleksandr Pryymak, Alex Rogers, Nicholas R. Jennings
Electronics and Computer Science
University of Southampton
Southampton, UK
{op08r, acr, nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper we present an approach for improving the accuracy of shared opinions in a large decentralised team. Specifically, our solution optimises the opinion sharing process in order to help the majority of agents to form the correct opinion about a state of a common subject of interest, given only few agents with noisy sensors in the large team. We build on existing research that has examined models of this opinion sharing problem and shown the existence of optimal parameters where incorrect opinions are filtered out during the sharing process. In order to exploit this collective behaviour in complex networks, we present a new decentralised algorithm that allows each agent to gradually regulate the importance of its neighbours' opinions (their social influence). This leads the system to the optimised state in which agents are most likely to filter incorrect opinions, and form a correct opinion regarding the subject of interest. Crucially, our algorithm is the first that does not introduce additional communication over the opinion sharing itself. Using it 80-90% of the agents form the correct opinion, in contrast to 60-75% with the existing message-passing algorithm DACOR proposed for this setting. Moreover, our solution is adaptive to the network topology and scales to thousands of agents. Finally, the use of our algorithm allows agents to significantly improve their accuracy even when deployed by only half of the team.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence*

General Terms

Algorithms, Performance, Reliability

Keywords

Self-organisation, Emergent behaviour, Distributed problem solving

1. INTRODUCTION

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

The problem of sharing information in large networked teams of hundreds and thousands of agents has recently received much attention in terms of how to facilitate the resolution of conflicting information and to improve its accuracy. In this paper we focus on a case when agents share *opinions* about the state of the common subject of interest, and these opinions may conflict. The aim of each agent is to maximise its *accuracy* by forming only *correct* opinions that correspond to the subject's true state. To fulfil its aim, the agent fuses opinions from other agents and forms its own opinion. However, in many decentralised systems, such as social communities and sensor networks, agents' interactions are restricted by a communication network. Thus, agents can receive opinions only from a limited number of their network neighbours. The complex topological properties of such communication networks [8] give rise to surprising and non-trivial *collective behaviour* in these opinion sharing processes [3, 10]. For example, a team may suddenly change its state when a large number of agents change their opinions in an *opinion cascade* after just a single new observation has been introduced [1]. Therefore, there is a crucial need to take into consideration and exploit the properties of collective behaviour in developing an agent-based approach for improving the accuracy of shared opinions.

Recently, Glinton, Scerri and Sycara [4, 5, 6] have presented an agent-based model of opinion sharing to analyse the impact of collective behaviour on the accuracy of the agents' opinions. In contrast to the classical models of opinion sharing [3], the researchers model observations of the common subject of interest by agents with noisy sensors. This approach enables us to reason about the accuracy of the opinions. Their analysis reveals that the accuracy dramatically increase in a narrow range of *social influence parameters* that encode how agents affect each other [4]. This narrow range correspond to a phase transition between a stable state of the team (where opinions are not shared) and an unstable one (where early and possible incorrect opinions are shared on a large scale). Close to this phase transition the number of agents that take part in an opinion cascade is distributed by a power law, and thus, opinion sharing exhibits *scale-invariant dynamics*. At this point, frequent small cascades prevent the team from overreacting to early and possibly incorrect opinions. While less frequent, large cascades share the locally supported opinions to the rest of the team. In Glinton et al.'s model the social influence parameters, on which the opinion sharing depends, are implemented as *importance levels* that each agent attributes to its neighbours' opinions. Unfortunately, it is impossible to predict the importance levels that introduce these scale-invariant dynamics since the properties of the communication network has a significant influence on the sharing processes and thus, analytical analysis cannot be applied to teams with complex communication networks [2]. In order to achieve the optimal parameter in such system, Glinton et al. proposed the Distributed Adaptive Communication for Overall Reliability (DACOR) algorithm [5]. DACOR is an online algorithm that adjusts the agents' importance levels according to the estimated local *branching factor* – the expected number of neighbours that would change their opinions following the change of an agent's opinion. In particular, it was found that in the area of optimal parameters the branching factor is close to 1.

However, actually performing a decentralised estimation of the branching factor by DACOR requires significant communication overhead compared to the opinion sharing itself. In many settings the capabilities of the agents are restricted and communication is limited to opinion sharing only. These restrictions can be found in many realistic settings, such as sensor networks where it is expensive to share data, or social communities where people rely on the opinions of others when they do not have enough resources or skills to analyse the original information themselves. Therefore, there is a need to address the open problem of improving the accuracy of shared opinions in settings where communication is strictly limited to opinion sharing. Moreover, to be applicable across a broad range of domains, a solution must adapt to the network topology. However, as our empirical evaluation reveals, the internal parameters of DACOR are very sensitive to the team's configuration and they have to be tuned individually for different domains.

To address these shortcomings, we present a decentralised algorithm for Adaptive Autonomous Tuning (AAT) of agents' importance levels. AAT improves the accuracy of the opinions in complex networks without introducing additional communication overhead. In contrast to DACOR, our algorithm relies solely on agents' local observations, rather than resource-intensive estimation of the branching factor. Our approach is based on the observation that opinions in the team becomes dramatically more accurate when the agents apply the minimal importance levels to their neighbours that still enable them to share opinions on the team scale. By meeting this condition at the individual agent level, AAT gradually tunes the team to the phase transition in the dynamics of the opinion sharing between the stable and the unstable state. In such settings the team exhibits significant improvement in the accuracy of agents' opinions since the team does not overreact to early and possibly incorrect opinions and the agents share opinions in smaller groups before a large cascade occurs.

In more detail, the contributions of this paper are:

1. We develop a novel decentralised algorithm, AAT, that improves the accuracy of the opinions in a large team with a complex communication network by exploiting the properties of its collective behaviour. Crucially, AAT is the first solution that operates when communication is strictly limited to opinion sharing, and is able to adapt to the specific communication network in which the agents find themselves.
2. We empirically evaluate AAT and show that it significantly outperforms the state-of-the-art solution, DACOR. Specifically, using AAT, 80-90% of the agents' typically form the correct opinion about the common subject of interest. This figure is significantly higher than 60-75% for DACOR, and close to 90-95% that can be reached by pre-tuning a team by an expensive empirical exploration of its parameters. Moreover, AAT introduces less computation expenses and each agent requires 10^4 times less actions than with DACOR.
3. We show that AAT is the first efficient solution designed to improve accuracy in teams with indifferent agents that do not participate in the optimisation process. Specifically, it significantly improves the accuracy with up to 50% of indifferent agents in the team. This implies that AAT potentially can be used by the large teams where it is impossible to update the behaviour of all agents, such as human-agent networks or heterogeneous sensor networks.

The remainder of this paper is organised as follows. In Section 2 the model of the environment, its properties and metrics are discussed. In Section 3 the agents' dynamics are analysed and AAT is presented. Then, in Section 4 AAT is empirically evaluated to demonstrate its efficiency in contrast to DACOR and it is compared with a team pretuned for the highest accuracy. Section 5 concludes this work.

2. PROBLEM DESCRIPTION

In this section, we formally describe an agent-based model of opinion sharing that was recently proposed and analysed by Glinton, Scerri and Sycara [4, 5, 6]. The aim of the model is to capture the complex dynamics of opinion sharing in a network of cooperative agents. In this model, some agents have access to noisy sensors, and they introduce to the team conflicting opinions of which only one is correct. Due to communication constraints agents can only share opinions with their network neighbours, without any additional information.

2.1 Model of Opinion Sharing

Formally, the Glinton, Scerri and Sycara model consists of a large set of agents $A = \{i^l : l \in 1 \dots N\}$, $N \gg 100$ connected by a undirected network $G(A, E)$ where E is the set of edges indicating which agents are neighbours and can therefore communicate. Each agent, $i \in A$ has a neighbourhood $D_i = \{j : \exists (i, j) \in E\}$ and the average number of neighbours is defined as the expected degree d, where $d = \sum_{i \in A} |D_i|/N$. We assume that the network is sparse $d \ll N$ in order to observe the cascading behaviour in the sharing process.

The aim of every agent, and eventually of the whole team, is to find the true state b of the common subject of interest, for example $B = \{\texttt{white}, \texttt{black}\}$, where $b \in B$. We support the assumption that B is binary following the argument that a binary choice can be applied to a wide range of real world situations [11]. However, our approach, presented later, does not rely on this limitation and can be extended for $|B| > 2$.

The goal of each agent is to form its own correct *opinion*, o_i, such that $o_i = b$. To recover this true state, agents rely on noisy sensors and their neighbours' opinions about the value of b. To decide which conflicting opinion to adopt, agent i forms its private belief $P_i(b{=}\texttt{white})$, which is the probability that $b = \texttt{white}$ (further denoted as P_i) and consequently $1 - P_i$ is the probability of $b = \texttt{black}$. The agent updates its belief starting from some initial prior P_i' and the ongoing belief is denoted by P_i^k where k is the current step of the belief update sequence.

Only a small subset of agents $S \subset A$, $|S| \ll N$ have noisy sensors and can make observations of the true state b. Each agent with a sensor $i \in S$ periodically receives an observation $s_i \in B$ with a low accuracy r ($0.5 < r \ll 1$), which is the probability of returning the true state b. To incorporate a new observation from the sensor into its belief, the agent uses formal reasoning based on Bayes' theorem:

$$P_i^k = \frac{c_{\text{upd}} P_i^{k-1}}{(1 - c_{\text{upd}})(1 - P_i^{k-1}) + c_{\text{upd}} P_i^{k-1}}, \tag{1}$$
$$\text{where} \begin{cases} c_{\text{upd}} = r & \text{if } s_i = \texttt{white} \\ c_{\text{upd}} = 1 - r & \text{if } s_i = \texttt{black} \end{cases}$$

After updating its belief with a number of observations the agent may become confident enough to form its own opinion o_i^k about the true state b. It does so once its belief P_i^k exceeds thresholds, following the opinion update rule:

$$o_i^k = \begin{cases} \texttt{undeter.}, & \text{initial, if } k{=}0 \\ \texttt{white}, & \text{if } P_i^k \geq \sigma \\ \texttt{black}, & \text{if } P_i^k \leq 1{-}\sigma \\ o_i^{k-1} & \text{otherwise} \end{cases} \tag{2}$$

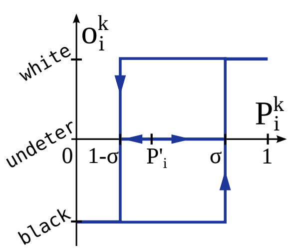

Figure 1: The opinion update rule

where thresholds $\{1{-}\sigma, \sigma\}$ are the *confidence bounds* and $0.5 < \sigma < 1$. The opinion update function has the shape of a sharp hysteresis loop (Fig. 1), and because sensors are noisy, it is possible that later observations will support the opposite opinion, and the agent may change its opinion.

Every time the agent changes its opinion, it communicates the new opinion to its neighbours. Consequently, these neighbours update their own beliefs and may form their own opinions. If the agent changes its opinion following a received opinion from its neighbour, it participates in an *opinion cascade* where a number of agents change their opinions in a sequence after a critical sensor observation. In order to incorporate opinions of the neighbours, the agent uses Bayes' theorem to update its belief similarity to sensor observations, such that when the agent receives new opinions from its neighbours $\{o_j : j \in D_i\}$, it uses the same belief update rule for each received opinion o_j:

$$\text{Eq. (1), where} \begin{cases} c_{\text{upd}} = t_i & \text{if } o_j = \texttt{white} \\ c_{\text{upd}} = 1 - t_i & \text{if } o_j = \texttt{black} \end{cases} \tag{3}$$

where $t_i \in [0, 1]$ is the *importance level*. This is the measure of the social influence of the neighbour's opinion (that is a conditional probability on opinions communicated from the neighbours). Note, the similarity with Equation 1 such that the importance level is analogous to the accuracy of a noisy sensor, r. However, unlike the accuracy r of a sensor, importance level t_i is unknown and each agent must find its value. In Section 3 we offer our algorithm for this purpose.

The agents in this model are cooperative and thus, they consider only the range $t_i \in [0.5, 1]$, where $t_i = 0.5$ indicates that the received opinion is ignored, and $t_i = 1$ is the maximum importance such that the agent changes its belief to $P_i^k = \{1, 0\}$ (depending on the received opinion) regardless of its previous value P_i^{k-1}. The model implies that the neighbours can be equally wrong in their opinions since sensor readings are introduced randomly. Therefore, it makes an additional assumption that the agent does not differentiate the sources of received opinions and applies the same importance level t_i for all its neighbours. We intend to relax this assumption in our future work and develop techniques that will help to make decisions about the importance of the opinions of each neighbours' individually.

Glinton et al. showed that this model exhibits emergent behaviour and the agents' opinions converge to the true state dramatically more often when the number of agents that take part in an opinion cascade is distributed by a power law, that is known as scale-invariant dynamics [5]. The importance levels are a key parameter which regulate the sharing process and thus, impact the distribution of sizes of opinion cascades. Unfortunately, it was shown that it is infeasible in the general case to predict the importance levels (t_{emr}), at which the emergent behaviour occurs, as this is highly dependent on topology of the network, the distribution of the priors of the agents' and the properties of the sensors. When the team operates with importance levels lower than the critical $\forall i \in A : t_i \ll t_{\text{emr}}$, it is in the stable state of its dynamics and the agents cannot form their own opinions because their beliefs never cross the confidence bounds. Conversely, the team is in the unstable state when $t_i \gg t_{\text{emr}}$, and the agents instantly form confident beliefs, propagate the first, possibly incorrect opinion, and do not benefit from the presence of multiple sensors in the team.

2.2 Performance Metrics of the Model

In order to measure the performance of the team, the model is simulated for a number of opinion dissemination rounds, $M = \{m_l : l \in 1 \dots |M|\}$, where in each round the new true state $b^m \in B$ is selected randomly. We observe the agents' final opinions, o_i^m, at the end of each round, m. Each round is limited by a large number of belief update steps, k, after which the team is likely to converge to the state where no agent is willing to change its opinion. The end of each round constitutes a certain deadline when the current true state expires. It may be followed by further rounds, in which case, the agents reset their beliefs and opinions to the initial values.

To measure the average accuracy of the agents' opinions at the end of each round, Glinton et al. [5] proposed a metric based on the accuracy of the team that was defined as the ratio between the number of dissemination rounds when the agents' final opinions are correct versus incorrect. This metric heavily penalises the team for disseminating incorrect opinions. However, it can be also maximised if a large proportion of the team does not form any opinion. This is somewhat problematic because we note that in many scenarios it is also important for the agents to form an opinion even if that opinion turn out to be incorrect. Thus, there is a need to balance both the need to be correct, and to actually form an opinion. Therefore we offer the accuracy metric that measures how often an agent forms the correct opinion on average:

$$R = \frac{1}{N|M|} \sum_{i \in A} |\{m \in M : o_i^m = b^m\}| \cdot 100\% \qquad (4)$$

Additionally we introduce a metric from a perspective of a single agent. Since it cannot determine when it has formed a correct opinion, the agent is interested to measure how often it forms an opinion. We denote this as an agent's *awareness rate*, h_i, that is the proportion of dissemination rounds where the agent i held an opinion rather than being undetermined compared to the total number of rounds:

$$h_i = \frac{|\{m \in M : o_i^m \neq \texttt{undeter.}\}|}{|M|} \qquad (5)$$

This myopic metric can be calculated locally by each agent and we use it as a basis of our algorithm later. Having introduced the model, we look next at algorithms which optimise the accuracy R, and in Section 4 we offer additional metrics to evaluate their efficiency.

3. AUTONOMOUS ADAPTIVE TUNING

In this section, we present our Autonomous Adaptive Tuning (AAT) algorithm, for improving the accuracy R of a complex communication network by exploiting its collective behaviour. In contrast to the existing algorithm, DACOR, our solution does not introduce communication overhead and communication is strictly limited to opinion sharing. Specifically, in order to estimate the local branching factor, DACOR requires that following a change of an agent's opinion, all its neighbours communicate on average d^2 additional service messages, where d is the expected number of neighbours.

We address this shortcoming by developing a new solution that updates agents' importance levels autonomously, relying on their local observations. Specifically, AAT is built on the observation that accuracy significantly increases when the dynamics of the opinion sharing is in the phase transition between the stable state (when opinions are not shared, $\forall i \in A : h_i \ll 1$) and an unstable one (when the first introduced opinion is propagated on a large scale, $h_i = 1$). This creates a condition where the team does not overreact to incorrect opinions and the agents share opinions in smaller groups before a large cascade occurs. To reach this area of optimal parameters, AAT gradually tunes an importance level of each agent individually.

The three stages of AAT are described in the following sections. Firstly each agent running AAT builds a set of candidate importance levels to reduce the search space for the following stages. Then the agent estimates the awareness rates of the candidate levels after each dissemination round. Finally, the agent selects an importance level to use in the following round, considering how close its estimated awareness is to the target awareness rate.

3.1 Candidate Importance Levels

In this section, we discuss how each agent running AAT selects a number of candidate importance levels, T_i, which reduces the continuous problem of selecting an importance level to use, t_i, from the range $[0.5, 1]$ to a discrete problem. In the general case T_i may be populated with importance levels drawn from the range $[0.5, 1]$ with a given step size, for example 0.01. However, by analysing the dynamics of an agent we can offer a solution for selecting a smaller number of candidate levels which will help AAT to converge to the optimal parameters faster.

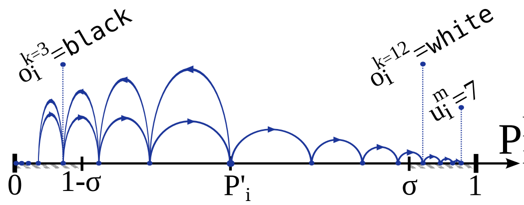

Figure 2: The sample dynamics of an agent's belief with marked steps when the agent changed its opinion.

Since the number of sensors is very small, we focus on the analysis of the agents without sensors who inform their beliefs using only their neighbours' opinions. For example, Figure 2 illustrates the sample dynamics of an agent's belief, P_i^k, where the agent i participated in 2 opinion cascades of conflicting opinions. Starting from its prior P_i', using Bayes' theorem the agent updates its belief with 4 neighbours' opinions that support '`black`' (i.e. 4 updates to the left from the prior P_i' that decrease the agent's belief $P_i^k(b = \texttt{white})$), after which the agent sequentially receives 11 opinions supporting '`white`' (i.e. updates to the right that increase P_i^k).

Clearly, the most important moments in this dynamic are the update steps when the agent change its opinion (steps $k = 3$ and $k = 12$) since only at these steps does the agent communicates a new opinion to its neighbours. To find all the cases whereby the agent can influence the local dynamics, we must find all the importance levels for which the opinion formation process may change. According to the opinion update rule (Eq. 2) we can limit this analysis only to those cases when the agent's belief coincides with one of the confidence bounds $P_i^k \in \{\sigma, 1{-}\sigma\}$. Considering also that the maximum number of opinions that the agent can receive is limited to the number of its neighbours, $|D_i|$, we can allow each agent to pre-calculate the candidate importance levels. Specifically, the agent has to find only those importance levels for which its belief coincides with one of the confidence bound $P_i^l \in \{\sigma, 1{-}\sigma\}$ in $l \in 1 \ldots |D_i|$ updates (see Eq. 3). By solving this problem, the agent constructs a set of the candidate importance levels that lead to opinion formation after receiving $1 \ldots |D_i|$ identical opinions and reaching the confidence bound σ or $1 - \sigma$:

$$T_i = \left\{ t_i^l : P_i^l(t_i^l) = \sigma,\ l \in 1...|D_i| \right\} \cup \left\{ t_i^l : P_i^l(1 - t_i^l) = 1 - \sigma,\ l \in 1...|D_i| \right\} \qquad (6)$$

As a result, the set of candidate levels is limited to twice the number of neighbours: $|T_i| = 2|D_i|$. This is a complete set of importance levels for which the agent forms an opinion on different update steps and it has to be initialised only once. Now, the agent has to form its preferences over these candidate levels to select the most appropriate one to use.

3.2 Estimation of the Awareness Rates

In this section we present the criteria according to which AAT selects an importance level from the candidates. As mentioned earlier, AAT is based on our observation that the accuracy, R, is maximised when the dynamics of opinion sharing is in a phase transition between stable and unstable state. In order to reach such optimal parameters the agents should use the minimal importance levels to their neighbours that still enable them to share opinions on the team scale.

The intuition is that in order to form an accurate opinion, the agent has to gather as many of its neighbours' opinions as possible before forming its own opinion. To do so, it has to use the minimal importance level from its candidate set. However, if all agents use the minimal importance level and wait until all their neighbours form opinions, a deadlock results where the opinion sharing stops. Therefore, each agent must apply a minimal importance level to the received opinions which guarantees that the agent actually forms its own opinion and shares it further.

In terms of the model we can formalise this, such that in order to maximise the accuracy, R each agent has to:

- Form its opinion, and thus, reach a high level of its awareness rate (h_i, the proportion of the rounds where the agent held an opinion rather being undetermined) since the agents with undetermined opinions decrease the team's accuracy;
- Form the correct opinion given its local view. Following the intuition above, in order to do so, the agent has to form an opinion as late as it is possible to gather the maximum number of neighbours' opinions.

To meet these conditions, the agent has to use the minimal importance level out of the candidates, $t_i^l \in T_i$, that always lead to an opinion formation ($h_i = 1$).

However, since sensors introduce observations randomly, the opinion sharing dynamic in the area of the phase transition exhibits stochastic behaviour. As a result, during some rounds opinions are not shared on a large scale and the agents' awareness rates suffer. Therefore, to improve the overall accuracy and find the exact position of the phase transition, each agent i has to compromise its own awareness rate, h_i. Specifically, the agent has to find the minimal importance level, t_i^l out of candidates T_i that delivers the *target awareness rate*, h_{trg}, that is slightly lower than the maximum, 1. Formally, each agent solves the following optimisation problem:

$$t_i = \underset{t_i^l \in T_i}{\arg\min} \, |h_i(t_i^l) - h_{\text{trg}}| \tag{7}$$

where $h_i(t_i^l)$ is the awareness rate that the agent achieves using importance level t_i^l. We analyse the impact of the specific value of h_{trg} on the accuracy in the empirical evaluation (Section 4.1).

Now, in order to perform the optimisation in Equation 7, the agent needs to calculate all awareness rates, $h(t_i^l)$, that would be achieved by using $t_i^l \in T_i$. However, according to the definition of the awareness rate, h_i (Eq. 5), it can be measured only for the importance level, t_i, that the agent currently uses. By analysing the process of the agents' belief update, we propose the following approach to estimate the awareness rate based on the local observation. Specifically, to estimate the awareness rate, $\hat{h}_i^l \approx h(t_i^l)$, the agent has to decide if its opinion could have been formed had it used an importance level, t_i^l, rather than the actually used t_i. We identify two cases that indicate this:

1. Consider the case that the agent used importance level t_i in round m and an opinion was formed, $o_i^m \neq$ undeter. According to the belief update function (Eq. 3) all higher importance levels, $t_i^l \geq t_i$, would have led to the more confident belief ($|P_i(t_i^l)| > |P_i(t_i)|$), and thus, to opinion formation.

Algorithm 1 AAT

Procedure UPDATE(i)
{Revises the current importance level after each round}
1: **if** OPINIONS RECIEVED : $u_i^m \neq 0$ **then**
2:     **for all** CANDIDATE LEVELS : $t_i^l \in T_i$ **do**
3:         **if** OPINIONFORMED$(t_i^l, t_i, m) =$ `True` **then**
4:             $\hat{h}_i^l =$ UPDATEAVERAGEAWARENESS$(\hat{h}_i^l,\ 1)$
5:         **else**
6:             $\hat{h}_i^l =$ UPDATEAVERAGEAWARENESS$(\hat{h}_i^l,\ 0)$
7:     $t_i =$ SELECTBYAWARENESS $\Big(\langle t_i^l,\ \hat{h}_i^l\rangle : l \in 1..|T_i|\Big)$

2. Otherwise, if the opinion was not formed, the agent can make a decision by comparing the number of updates it has observed and the number required for the candidate level, t_i^l, to form an opinion. Specifically, the minimal number of belief updates required to form the opinion with the candidate level, t_i^l, can be calculated by recursively updating the agent's belief (see Eq. 3) starting its prior until it exceeds one of the confidence bounds: σ for updates with t_i^l, or $1-\sigma$ with $1-t_i^l$. We denote this function as $u(t_i^l, P_i', \sigma)$. At the same time, during the dissemination round the agent can observe the maximum number of updates it has made in favour of any conflicting opinion starting its prior. We denote this value as u_i^m. In Figure 2 it is observed on the last belief update step $u_i^m = |4-11| = 7$. Finally, the opinion should have been formed when the the number of updates required for the candidate t_i^l is smaller or equal than the observed number u_i^m.

Combining these cases, we construct a boolean function that returns `True` if the agent might have formed an opinion in the current round, m using importance level t_i^l with actual importance level t_i:

$$\text{OPINIONFORMED}(t_i^l, t_i, m) = \Big(o_i^m \neq \texttt{undeter.} \wedge t_i^l \geq t_i\Big) \vee u_i^m \geq u(t_i^l, P_i', \sigma) \tag{8}$$

Following the definition of the awareness rate (Eq. 5), to estimate the awareness rates for the candidate levels the agent has to measure the proportion of dissemination rounds $m \in M$ for which the condition above was matched:

$$\hat{h}_i^l = \frac{|\{m \in M : \text{OPINIONFORMED}(t_i^l, t_i, m) = \texttt{True}\}|}{|M|} \tag{9}$$

Algorithm 1 describes the core procedure of AAT that implements this approach to estimate awareness rates and is executed after each dissemination round. If no opinions were received ($u_i^m = 0$), the agent cannot form its own opinion with any of the importance level, and thus this case is limited by the condition on line 1. In lines 2-6, AAT updates the estimates of the awareness rate for each of the candidate levels according to the procedure described above. Now, according to optimisation problem the agent solves (Eq. 7), it has to select the importance level (line 7) that delivers the awareness rate closest to the target, h_{trg}, considering the high interdependence between agents' choices.

3.3 Strategy to Select an Importance Level

The agents' opinions are highly interdependent and an importance level chosen by a single agent eventually affects the dynamics and awareness rates of all agents. Therefore, if the agent would greedily select its importance level according to the definition of its optimisation problem (Eq. 7), it may dramatically change local dynamics. Instead, the agent has to employ a strategy with less dramatic changes in its dynamics, in order that entire team estimate awareness rates more accurately and converge to the solution faster.

To construct such a strategy, we note that since the lowest importance level, t_i^1, from the candidates in ascending order, requires more sequential updates to cross one of the confidence bounds, while the largest $t_i^{\max}$ requires less, then the awareness rates are distributed as a hill with a peak for the largest importance level, $t_i^{\max}$. Therefore, we offer a *hill-climbing strategy* that makes use of this observation. If the awareness rate delivered by the currently used importance level, $t_i = t_i^l$, is lower than target $\hat{h}_i^l < h_{\text{trg}}$, the agent must increase the importance level to the closest larger one (i.e. $l = l + 1$). Conversely, if the closest lower importance level is estimated to deliver an awareness rate higher than $\hat{h}_i^{l-1} > h_{\text{trg}}$, the agent chooses to use it in the next round (i.e. $l = l - 1$). Our empirical evaluation confirmed that the hill-climbing strategy delivers the higher accuracy then the greedy strategy and for brevity we present results only of AAT based on it.

4. EMPIRICAL EVALUATION

To empirically evaluate the performance of AAT and the existing DACOR, we consider a wide range of parameters in order to examine their adaptivity and scalability. Specifically, we evaluate the accuracy of teams with $N \in \{150 \ldots 2000\}$ agents on networks with a variable expected degree, $d \in \{4 \ldots 12\}$. The maximum size of a team is limited due to the high computational expenses required to empirically pretune a team for the highest accuracy that we use later as a benchmark. We consider the following network topologies widely used in the literature: (a) a connected random network; (b) a scale-free network with clustering factor $p_{\text{cluster}} = 0.7$ [7]; (c) a small-world ring network with $p_{\text{rewire}} = 0.12$ of randomised connections. [9]. New opinions are introduced through a small number of sensors ($|S| = 0.05N$ with accuracy $r = 0.55$) that are randomly distributed across the team. To simulate a gradual introduction of new opinions, only 10% of sensors make new observations after the preceding opinion cascade has stopped. Finally, all agents are initialised with the same confidence bound $\sigma = 0.8$, initial opinion $o_i^0 = \texttt{undeter.}$, and individually assigned priors P_i' that are drawn from a normal distribution $\mathcal{N}(\mu = 0.5, s = 0.1)$ within the range of the confidence bounds $(1 - \sigma, \sigma)$.

Before every round m we randomly choose the true state $b^m \in B$. Each round stops after 3000 sensors' observations and sequential opinion cascades. After this number of observations, the opinions of the agents with sensors converge to the true state, and thus, the sharing process stops. The end of each round constitutes a deadline when the current true state expires, and agents reset their beliefs and opinions to the initial values. AAT and DACOR tune the importance levels in the first 150 rounds, then the metrics are measured over the following 150 rounds. Error bars in figures indicate the standard errors across a designated number of network instances in each case.

4.1 Selection of the Target Awareness Rate

We first analyse the performance of our algorithm AAT with a regard to its single parameter – the target awareness rate h_{trg}. The analysis supports our earlier assertion, that h_{trg} has to be slightly lower than 1 to tune the team to the area of optimal parameters. Figure 3 shows that the highest accuracy achieved when $h_{\text{trg}} = 0.9$ regardless of the topology of the network. The accuracy significantly drops for the higher values of h_{trg} (Fig. 3a) since agents select much larger importance levels (Fig. 3b) to form opinions out of smaller number of observations. Thus, they become overconfident and the whole team converges to the early opinion without fusing it with later observations that might be more accurate. Considering the results, in our further evaluation we use $h_{\text{trg}} = 0.9$.

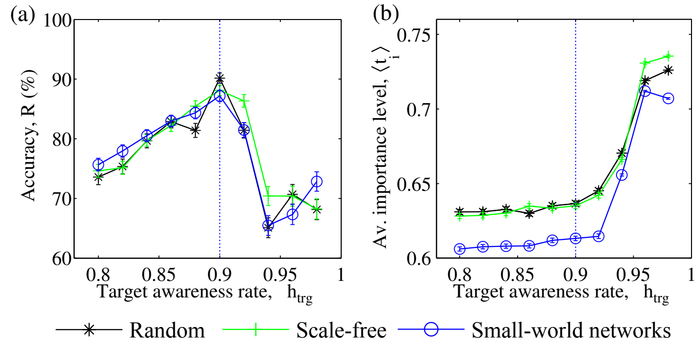

Figure 3: (a) The accuracy and (b) the average importance level achieved by AAT, both depend on the target awareness rate h_{trg} (40 instances of each topology with $N = 1000$ and $d \in \{4 \ldots 12\}$).

4.2 Accuracy of the Opinions

We now benchmark AAT against three alternative solutions. First, we compare against DACOR (with parameters $uA = 10, \gamma = 0.001, \beta = 0.1$ selected to maximise the accuracy of a random network with $d = 8$), the current state of the art solution in this setting. In addition, we also benchmark against a team pre-tuned for the highest performance on a specific network instance. In more detail, to pre-tune a team, we perform a resource intensive empirical exploration of each network instance with fixed importance levels $\forall i \in A : t_i = t$, where $t \in (0.5, 1)$ with a step of 0.05 over $|M| = 150$ rounds. Then we choose the importance level t_{emr} at which the team exhibits the highest accuracy. Note, that this is not the optimal solution, as it is infeasible to explore the whole domain where agents may have different importance levels. Still, this approximation exhibits a high accuracy of 90 to 97% and shows its level that can be achieved by fine tuning. However, t_{emr} varies between different network instances since the area of optimal parameters is very narrow and dependent on the team's configuration. Therefore, to illustrate how difficult such tuning is in practice, we also benchmark against the team with the average $\langle t_{\text{emr}} \rangle$ for the networks of the same size and topology.

The results of the accuracy benchmark are shown in Figure 4a. As can be seen, AAT shows accuracy close to the results of the pre-tuned teams and significantly outperforms the existing solution, DACOR, for all network topologies. AAT scales well, since it reaches the stable accuracy around of 86 to 88% for teams larger than 1000 agents. However, it declines as the team size becomes lower than 1000 agents. This is due to the fact that the properties of collective behaviour are less distinct in smaller teams. Analysis of the

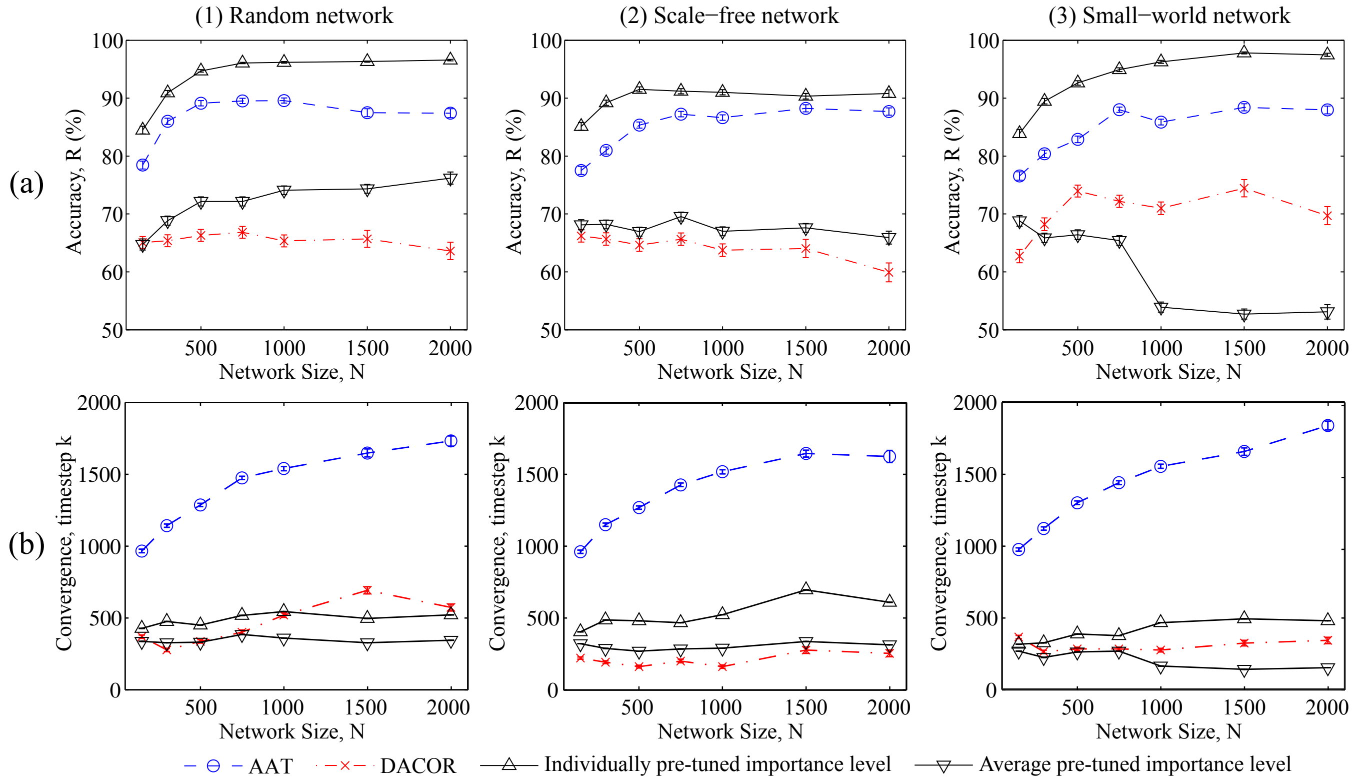

Figure 4: (a) The accuracy and (b) the convergence of a team with AAT, DACOR, and pre-tuned importance levels (40 instances of a each topology and network size with $d \in \{4 \dots 12\}$).

results also show that DACOR, unlike our adaptive AAT approach, is highly dependent on parameters which have to be individually tuned for specific domains. In most cases DACOR has tuned the team to the unstable state where early and possibly incorrect opinions are shared on a large scale. Finally, the low accuracy achieved by teams with $\langle t_{\text{emr}} \rangle$ indicates a clear need for an algorithm such as AAT that can efficiently tune each team individually.

4.3 Opinion Convergence

AAT tunes the team into a phase transition between stable and unstable states where the sharing processes are the slowest that still enable all agents to form their opinions. However, this also implies that agents with AAT may form their opinion slower. To measure the timeliness of the opinion, we offer a *convergence metric* that is the average number of timesteps required for a team to reach the accuracy of $R \geq 80\%$. In order to avoid distortion of its average value, we exclude dissemination rounds when the team did not reach the threshold level of the accuracy.

The results shown on Figure 4b indicate that convergence time for AAT growth steadily with the size of a team. This fact can explained by the increasing sparseness of the network since its degree d is fixed. This results in a slower sharing process as the shortest path increases as well.

DACOR exhibits much faster convergence since it tunes the team into the area of the unstable state. This also explains its low accuracy discussed above. By contrast, most of the teams with the average of pre-tuned levels exhibit stable dynamics and do not share opinions on a large scale, while some are in unstable state that result in a fast convergence.

Finally, the individually pre-tuned teams exhibit relatively fast convergence. This indicates a drawback of AAT that has to be addressed. Specifically, the online approach used to build AAT results in slower convergence and alternative solutions, such as offline pre-tuning, may exhibit equal or higher level of accuracy with significantly faster convergence at the same time.

4.4 Communication and Computation Expenses

AAT is designed to improve accuracy without introducing additional communication over opinion sharing. We compare in Figure 5a the number of messages that agents exchange while the team is tuned by AAT, DACOR, and the minimal number of messages required to share an opinion on a team scale in a single cascade. The latter represents the *minimal communication*, when agents share their opinions only once to the neighbourhood, and thus, communicate in total dN messages. The average number of messages for a team with AAT is similar to the minimal communication, since during some rounds a team does not disseminate opinions on a large scale (as the result of $h_{\text{trg}} < 1$).

In addition, AAT requires radically less actions by the agents that are the changes of the importance levels than DACOR in the process of tuning. AAT updates an importance level only once at the end of each round, while DACOR updates an agent's importance level if any of its neighbours has received new opinion. The results that represent computational expenses, are shown in Figure 5b. Both metrics show that AAT is a highly scalable solution.

4.5 Team with Indifferent Agents

Finally, AAT is robust to the presence of the agents that are *indifferent* and do not participate in the optimisation process. This is due to the fact that the agents running AAT

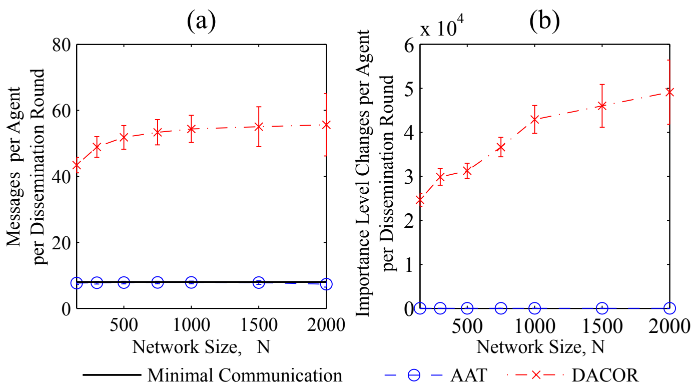

Figure 5: (a) Number of messages and (b) number of importance level changes, for an agent per dissemination round (averaged over all experiments shown in Figure 4).

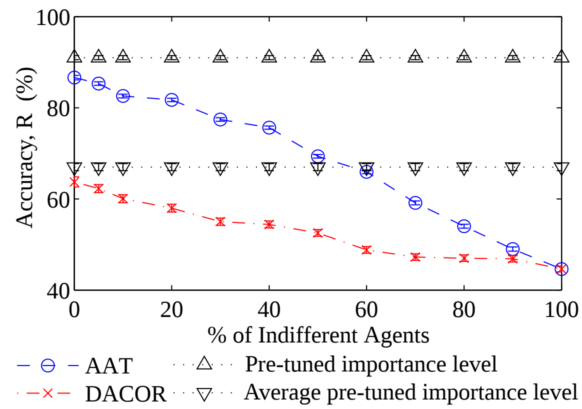

Figure 6: The accuracy of the team with indifferent agents (40 instances of a scale-free network with $N = 1000$, $d \in \{4 \dots 12\}$).

tune their importance levels autonomously by adapting to their neighbourhood. Thus, they mitigate the negative effect introduced by the indifferent agents.

We illustrate this by evaluating a team with a variable number of indifferent agents that are randomly distributed across its population. The importance levels of indifferent agents are not dynamically determined by AAT or DACOR algorithms, but fixed and uniformly selected from the range close to the critical importance level $[0.55, 0.75]$. The results in Figure 6 shows that AAT with up to 50% of indifferent agents delivers higher accuracy than can be achieved by using $\langle t_{\text{emr}} \rangle$. This shows the direct benefit from deploying AAT even on half of the agents in a team over predicting the critical importance level by analysing a number of similar teams. Similar results are obtained for the other topologies and team sizes.

5. CONCLUSIONS

In this paper, we developed a novel decentralised algorithm, AAT, which significantly improves the accuracy of agents' opinions by exploiting the properties of collective behaviour in large networked teams. This is the first solution that can be used by teams with complex communication networks in the settings when communication is strictly limited to opinion sharing. We showed that AAT significantly outperforms the existing algorithm, DACOR, that also introduces additional communication to operate and requires higher computational cost. The accuracy exhibited by AAT is close to the highest accuracy that can be achieved by individually pre-tuning a team by the resource expensive empirical exploration of its parameters. Moreover, we showed that AAT is scalable, adaptive to the team's configuration and robust to the presence of indifferent agents that do not participate in the optimisation process. Finally, since AAT relies only on local view, the importance levels it estimates can be used as the initial trust levels for an elaborate trust models when no additional information is available.

Our future work in this area is to relax an assumption that the agents do not differentiate between their neighbours. This will require a new algorithm that estimates individual importance levels for each neighbour based on their opinion dynamics. Additionally, we also intend to address an outlined problem of developing an attack resistant solution that will help to mitigate the negative influence of malicious agents [6].

6. REFERENCES

[1] S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.

[2] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, and D. Hwang. Complex networks: Structure and dynamics. *Physics Reports*, 424(4-5):175–308, 2006.

[3] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2):591–646, May 2009.

[4] R. Glinton, P. Scerri, and K. Sycara. Towards the understanding of information dynamics in large scale networked systems. In *Proceedings of 12th International Conference on Information Fusion (FUSION'09)*, pages 794–801, Seattle, Washington, USA, 2009.

[5] R. Glinton, P. Scerri, and K. Sycara. Exploiting scale invariant dynamics for efficient information propagation in large teams. In *Proceedings of 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'10)*, pages 21–28, Toronto, Canada, 2010.

[6] R. Glinton, P. Scerri, and K. Sycara. An Investigation of the Vulnerabilities of Scale Invariant Dynamics in Large Teams. In *Proceedings of 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'11)*, pages 677–684, Taipei, Taiwan, 2011.

[7] P. Holme and B. Kim. Growing scale-free networks with tunable clustering. *Physical Review E*, 65(2):2–5, Jan. 2002.

[8] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney. Statistical properties of community structure in large social and information networks. *In Proceeding of the 17th International Conference on World Wide Web (WWW'08)*, pages 695–705, 2008.

[9] M. Newman. Renormalization group analysis of the small-world network model. *Physics Letters A*, 263(4-6):341–346, Dec. 1999.

[10] D. J. Watts. A simple model of global cascades on random networks. *In Proceedings of the National Academy of Sciences of the United States of America*, 99(9):5766–5771, Apr. 2002.

[11] D. J. Watts and P. S. Dodds. Influentials, networks, and public opinion formation. *Journal of Consumer Research*, 34(4):441–458, 2007.