

Speed-vs-Accuracy Tradeoff in Collective Estimation: An Adaptive Exploration-Exploitation Case*

Mohsen Raoufi^{1,2}, Heiko Hamann³ and Pawel Romanczuk^{1,4}

Abstract—The tradeoff between accuracy and speed is considered fundamental to individual and collective decision-making. In this paper, we focus on collective estimation as an example of collective decision-making. The task is to estimate the average scalar intensity of a desired feature in the environment. The solution we propose consists of exploration and exploitation phases, where the switching time is a factor dictating the balance between the two phases. By decomposing the total accuracy into bias and variance, we explain that diversity and social interactions could promote accuracy of the collective decision. We also show how the exploration-vs-exploitation tradeoff relates to the speed-vs-accuracy tradeoff. One significant finding of our work is that there is an optimal duration for exploration to compromise between speed and accuracy. This duration cannot be determined offline for an unknown environment. Hence, we propose an adaptive, distributed mechanism enabling individual agents to decide in a decentralized manner when to switch. Moreover, the spatial consequence of the exploitation phase is an emergent collective movement, leading to the aggregation of the collective at the iso-contours of the mean intensity of the environmental field in the spatial domain. Examples of potential applications for such a fully distributed collective estimation model are spillage capturing and source localization.

I. INTRODUCTION

The ability of a collective to estimate a quantity, called collective estimation, has been considered as an example of collective decision-making (CDM) [1]. The idea, that the average of many imperfect estimations can, under appropriate conditions, be surprisingly accurate (even more accurate than the estimation of individual experts) is the so-called wisdom of crowds effect [2]–[4]. Researchers have shown that there are some conditions needed to be satisfied in order to have a “wise” crowd, and in the case of violation of such conditions the collective intelligence may instead turn to “madness” [5]–[7]. Although it is still an open question whether there is any sufficient condition to make collectives intelligent, some features, such as diversity, have been proposed as a promoter to the wisdom of crowds [4], [8].

*This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2002/1 “Science of Intelligence” – project number 390523135.

¹ Mohsen Raoufi and Pawel Romanczuk are with Institute for Theoretical Biology, Department of Biology, Humboldt Universität zu Berlin, Berlin, Germany

² Mohsen Raoufi is with Department of Electrical Engineering and Computer Science, Technical University of Berlin, Berlin, Germany mohsen.raoufi@hu-berlin.de

³ Heiko Hamann is with Institute of Computer Engineering, University of Lübeck, Lübeck, Germany hamann@iti.uni-luebeck.de

⁴ Pawel Romanczuk is with Bernstein Center for Computational Neuroscience, Berlin, Germany pawel.romanczuk@hu-berlin.de

In particular, the “diversity prediction theorem” states that the heterogeneity of decision-makers makes the aggregated estimation more accurate [9]. It is also known that, if the estimation errors of individuals are *unbiased* (yet even large), averaging the individuals estimations can cancel the errors out and promote collective accuracy [10], [11]. On the other hand, the effect of social influence, as a means to decrease diversity, has been controversially discussed in the wisdom of crowds context [10]–[12]. In this paper, we specifically focus on the diversity of information available within the collective prior to the social interaction, and show how the aggregation, as a consensus mechanism, decreases the total estimation error while keeping the collective estimation unbiased.

An increase in the accuracy of a decision-maker, whether individually or collectively, typically comes at the expense of a slower decision speed, which forms the well-known speed-vs-accuracy tradeoff (SAT) [13]. This ubiquitous phenomenon has been studied in a variety of systems, ranging from animal decision-making (DM) [14], [15], to human CDM [16], to neural systems [17]. Various factors modulate the SAT; in natural systems, it has been shown that the ability to modify individual responsiveness can have strong impact on the group behavior, especially under predation risk or environmental uncertainty [18], [19]. Similarly, in the swarm robotic systems, the effect of different parameters, such as neighborhood size, on the SAT has been studied [20], [21]. In another study [22], the SAT was modulated by algorithm parameters in CDM of 100 simulated Kilobots with the chief objective to make binary decisions about a feature in the environment. Changing the time for switching from exploration to exploitation influenced the accuracy of CDM for a swarm of simulated Kilobots in a best-of- n problem [23].

Any system with the ultimate goal of searching must handle the tradeoff between exploration and exploitation [24]. The exploration-exploitation tradeoff has been largely investigated in different CDM scenarios [23], including multi-armed bandits [25], [26]. Although a clear definition of exploration is context-dependent, in a spatial search for a resource, one can identify “pure” exploration phase corresponding to the accumulation of information about the environment. It has been shown that in this case, the exploration-exploitation tradeoff is analogous to the SAT [27]. We refer to exploitation as the ability of agents to interact with each other to exploit the available information within the collective. On the contrary, the exploratory behavior is considered as any action that can possibly provide new information to the collective from the environment and thus, increase

the diversity of the available information. This is in line with the definition that “exploration provides the agents with the opportunity for obtaining information” [27]. Considering spatial search for information in this paper, exploration has the same meaning as diffusion of agents here. We would like to note that diffusion typically refers to the stochastic motion of individual particles/agents and both in physical and biological sciences typically assumes independence between different particles (see e.g. [28] and [29]).

The organization of the paper is as follows: in section II we first define the problem and propose our exploration-exploitation solution. In the simplest case, the switching time from exploration to exploitation is a fixed parameter. We also propose an adaptive mechanism for switching, where each agent decides based on a simple experience-based heuristic when to switch from exploration to the exploitation phase. In section III, we introduce metrics to evaluate the collective estimation performance and discuss the experimental setups. The results are presented in section IV, before the conclusion and future outlook.

II. METHOD

A. Problem Definition

We address the collective estimation of a group of agents in an exploration-exploitation task. We assume that a collective of N agents is meant to search the environment and gather information of a specific, measurable, continuous feature. The intensity of the feature is distributed in the environment, and based on which the agents exchange information and the collective arrives at a decision about the average intensity within the specific, bounded region. The cardinal objective of the task is to collectively decide whether to take a specific action, based on the estimated intensity. We also show that the proposed generic distributed process maps to a concrete application to collectively identify and self-organize towards (mean) isolines of a spatial variable feature field, which may be of direct relevance for real-world applications. Such a collective motion happens as a result of the tendency of agents to move toward the locations of consensus by incorporating only local social information based on limited individual perception.

The exploration-exploitation task is decomposed into two distinct phases, and in the following, we explain how agents perform the sub-tasks in each phase. We also propose an adaptive, decentralized mechanism for the agents to decide when to switch from exploration to exploitation. The flow-chart in Figure 2 depicts the whole method including the proposed adaptive mechanism for switching to exploitation.

B. Exploration Phase

In the first phase of the task, agents explore the environment to reach uncovered areas in the arena by doing random diffusion. In this phase, agents do their task completely independent, and there is no need for information exchange between agents, i.e. agents do “pure” exploration behavior in this phase. Among the variety of potential exploration

methods, we modeled the random movement of agents as a *random walk*, where an agent i at time step t turns a small random increment and takes a step along its orientation ψ_i^t , according to

$$\psi_i^{t+1} = \psi_i^t + r_\psi \zeta_i^t, \quad (1)$$

where, ζ_i^t is a random number drawn from a uniform distribution in the range $[-\pi, \pi]$, and r_ψ is a parameter defining the rate of random orientation change. Once agents update their orientation, they take a step of the fixed size λ . The equation of motion for the update of the position vector of agent i at time step t reads:

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \lambda \begin{bmatrix} \cos \psi_i^t \\ \sin \psi_i^t \end{bmatrix}. \quad (2)$$

Agents continue to explore until they individually decide to switch to exploitation phase. This decision can be made by setting a fixed predefined switching time or, as we propose later, using an adaptive mechanism in an online manner. Therefore, the exploration phase only changes the spatial distribution of the agents in the environment at the onset of exploitation. We designed this model in order to specifically investigate the role of information diversity for collective estimation tasks, and how it can be influenced. In our model, we consider a finite, rectangular environment with reflecting boundary conditions. If an agent hits the wall it turns randomly and continues the exploration. The fact that agents do random walk enables them to eventually turn and leave the wall even in the case that they are not able to detect the collision.

C. Exploitation Phase

The aim of the exploitation phase is to aggregate or to combine the information distributed within the collective [30].

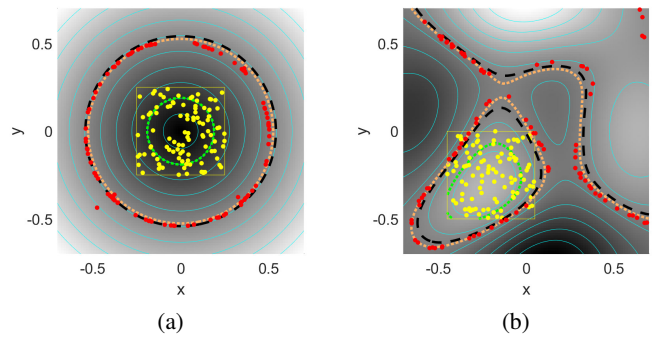


Fig. 1: Snapshot of the arena showing the distribution of the feature intensity in the environment. The yellow and red points are the initial and final position of the agents, the green and orange dotted lines are the initial and final intensity contours of the collective mean, respectively. The yellow rectangle depicts the initial placement patch of agents. The cyan curves show iso-contours of the intensity at different levels, and the black dashed one is the environment average (ground truth); a: cone-shape, uni-mode function, b: multiple peaks, multi-modal function.

The aggregation can be defined in two distinct domains, the information space and the spatial (physical) domain. In order to enable the agents to aggregate in both domains we considered two different mechanisms. The first mechanism is social interaction with local neighboring agents, which promotes the aggregation of information, and the second one is the motion of agents in the spatial environment in response to the social information. Although the second mechanism directly causes spatial aggregation, it can indirectly promote information aggregation as well. During the exploitation phase, agents constantly measure the intensity of the desired feature in the environment. We assume noisy measurements which is modeled via the following equation:

$$z_{s,i}^t = g(\mathbf{x}_i^t, t) + \sigma \mathcal{N}(0, 1), \quad (3)$$

where agent i senses the intensity as $z_{s,i}^t$, and $g(\mathbf{x}, t)$ is the spatial distribution function of the intensity on point \mathbf{x} at time-step t . The second term represents the uncertainties in perception, where $\mathcal{N}(0, 1)$ is a standard normal (white) noise and σ is the noise coefficient.

The social interaction happening in this phase promotes the wisdom of crowds effect [4], [31], by enabling the agents to average their imperfect estimates of environmental cues [24]. Agents locally exchange information if their neighbors are close enough. The communication range d defines neighborhood for agents. For a focal agent i , N_i refers to the number of other agents j in its vicinity with $|\mathbf{x}_i^t - \mathbf{x}_j^t| < d$. We assume a restricted spatial region of interest modeled by a rectangular arena with reflecting boundary conditions, and a finite time budget for the collective decision task. Note, that the consideration of a finite arena minimizes the chance of the interaction network to collapse or fragment. In open space, with increasing exploration time, a random exploration behavior results in more and more agents getting dispersed too far and getting disconnected, with a low probability of getting reconnected. As a consequence, these agents cannot participate in collective exploitation. Even in a restricted arena, there is a finite probability to observe disconnected clusters of agents during the experiment. This is a consequence of the minimal model of independent random exploration without any control on the connectivity of the network.

Once an agent switches to the exploitation phase, it starts to interact with local neighbors who are also in the exploitation phase, by exchanging information and accordingly update its state (or opinion) about the environment. By repeating these updates, agents gather more information from the collective, and use their *a priori* data (the raw data that they measured from the environment) to update their *posterior* information. This updating of the opinion signal is modeled as a weighted average of an individual opinion with memory and a collective signal. To do so, at time step t , a focal agent i collects the current local decision state $z_{col,i}^t$ by averaging the votes of itself and its neighbors (Eq. 5), which for a binary state space would result in a majority rule. Agent i is able to communicate only with N_i neighbors within its communication range d .

Then, agent i incorporates this instantaneous local collective signal (collective signal, in short) $z_{col,i}^t$ and its own memory $z_{m,i}^t$ using a weighting factor α to update its memory. This interaction is inspired by the DeGroot model [32]. By observing the opinion of neighbors, agents revise their own opinion in a way that their beliefs become more similar to their social neighborhood [11]:

$$z_{m,i}^{t+1} = \alpha z_{m,i}^t + (1 - \alpha) z_{col,i}^t. \quad (4)$$

The collective signal is calculated using the following voting method:

$$z_{col,i}^t = \frac{z_{s,i}^t + \sum_{j \in N_i} z_{m,j}^t}{1 + N_i}. \quad (5)$$

By putting this signal into the base equation 4, three different components show up, meaning that for an agent to make a decision, three sources of information can be taken into the equation: individual memory, individual perception, and social impact. We obtain

$$z_{m,i}^{t+1} = \alpha z_{m,i}^t + \frac{1 - \alpha}{1 + N_i} z_{s,i}^t + \frac{1 - \alpha}{1 + N_i} \sum_{j \in N_i} z_{m,j}^t. \quad (6)$$

The weighting between the three components is tuned by α , and is influenced by the network structure. One can generalize the model by increasing the degrees of freedom to two, by adding an additional parameter, whereby the sum of the weights of the three sources of information must equal to one.

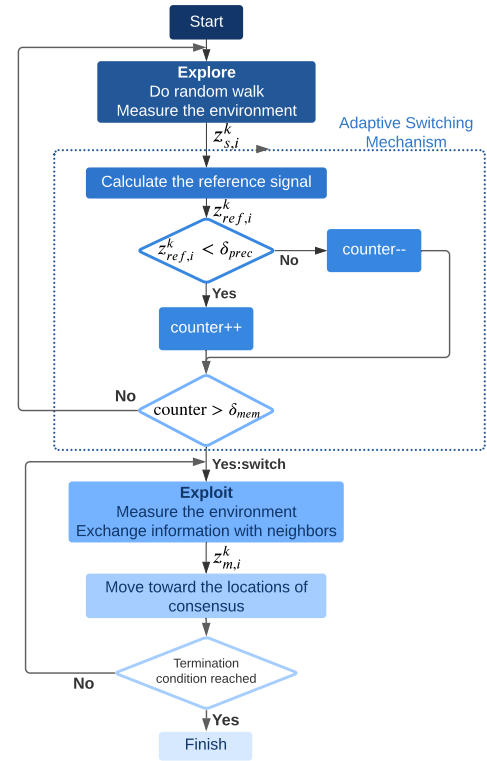


Fig. 2: The flow-chart of the exploration-exploitation method including the adaptive switching mechanism.

The second type of aggregation in physical space results from individual agents moving towards others with sharing a similar opinion (i.e. like-minded) about the environment, the so-called assortative mixing, or homophily. To implement such social movement, we define the objective function f_i^t , which is simply the difference between the current sensed signal and the collective signal for agent i :

$$f_i^t = \frac{1}{2}(z_{s,i}^t - z_{col,i}^t)^2. \quad (7)$$

The minimum of the objective function lies where an agent finds all of its neighbors signaling the same value as itself. In other words, an agent assumes the true value to be the collective signal, and tries to move to a location, where the sensed signal matches best the collective signal. To do so, agents need to find a (sub-)optimal step to take and update their position according to:

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \lambda_i^t. \quad (8)$$

There are various approaches to determine the step vector λ_i^t that optimizes the corresponding objective function. A classical approach is to apply a numerical pseudo-gradient descent method. Throughout the paper, we assume that agents have only access to local information, either from their neighbors, or the environment at their position. Thus, the pseudo-gradient methods that we propose do not involve any absolute, global position of agents in the environment, but only the last step that they took. In the following, two variants of such a method are described.

1) *Basic Pseudo-Gradient Descent Method:* Agents directly evaluate the objective function by sampling the environment while they are taking steps. The differentiation of the objective function over the step, that an agent took in the last time step, defines the slope of the objective function at the position vector $\mathbf{x}_i^t = [x_i^t, y_i^t]^T$. An approximation of the gradient is achieved by using a decaying memory of this differentiation:

$$\nabla_{\mathbf{x}} f_i^t = \beta \nabla_{\mathbf{x}} f_i^{t-1} + (1 - \beta) \begin{bmatrix} \frac{\partial f_i^t}{\partial x_i^t} \\ \frac{\partial f_i^t}{\partial y_i^t} \end{bmatrix}, \quad (9)$$

$$\frac{\partial f_i^t}{\partial x_i^t} \approx \frac{\Delta f_i^t}{x_i^t - x_i^{t-1}}, \quad \frac{\partial f_i^t}{\partial y_i^t} \approx \frac{\Delta f_i^t}{y_i^t - y_i^{t-1}}. \quad (10)$$

A random walk component is added with a weighting factor β to the gradient of the objective function. Such random components prevent the optimization to be greedy and compensate for the exploration that still is needed for optimizing the objective function. Overall, the step vector is determined with the fixed size λ as follows:

$$\lambda_i^t = \lambda \left(-(1 - r_\lambda) \frac{\nabla_{\mathbf{x}} f_i^t}{|\nabla_{\mathbf{x}} f_i^t|} + r_\lambda \eta_i^t \right), \quad (11)$$

in which, r_λ and η_i^t are the weight of random walk in the gradient descent, and a random variable drawn from a uniform distribution in the range $[-1, +1]$, respectively.

There are various sources contributing to the variability of the gradient approximation. For instance, an agent in

a fixed spatial position, perceives a time-varying collective signal as the neighboring agents are constantly updating their estimations. In addition, the gradient is a spatial function because not only the intensity function itself is a spatial function, but also the neighbor set changes due to the motion of the focal agent. So, as an alternative to the pseudo-gradient descent method mentioned above, we propose another method for approximation of the gradient that considers these factors explicitly and makes the approximation more reliable. Instead of directly approximating the gradient of the objective function, it approximates the gradient of the intensity function.

2) *Extended Pseudo-Gradient Descent Method:* An alternative of the previous gradient descent method for the same objective function can be derived by applying some calculus:

$$\nabla_{\mathbf{x}} f_i^t \approx \frac{N_i}{N_i + 1} (z_{s,i}^t - z_{col,i}^t) \nabla_{\mathbf{x}} z_{s,i}^t, \quad (12)$$

which means that each agent compares its sensed value to the decision of its neighbors. According to the sign of this difference decides to go either upward or downward on the gradient of the intensity function.

D. Adaptive Switching Mechanism

For an agent to know when to finish the exploration and switch to the exploitation phase requires a priori information about the environment, which in most of realistic applications is not available. In this part, we propose an adaptive, distributed method that enables agents to decide when to switch by measuring the quality of their exploration. In other words, the proposed mechanism quantifies the information gain of the input, and provides an index for individual switching. In order to explain the mechanism, we need to define signals determining the exploration index.

The first signal is the average of lower and higher bounds of the measured intensity that each agent has experienced during the exploration phase. Agents need to keep track of the min and max values of the measured signal to calculate their average value $z_{avg,i}^t$ as:

$$z_{avg,i}^t = \frac{1}{2} (\min\{z_{s,i}^l | l \in \{1, 2, \dots, k\}\} + \max\{z_{s,i}^l | l \in \{1, 2, \dots, k\}\}) \quad (13)$$

The second signal is a lag transformation of the average signal, or alternatively called the exponentially decaying average with factor β_{lag} . Applying this operator on a signal adds inertia, which makes it reluctant to change. Thus, the resulting signal becomes more stable and converges to the steady-state value more slowly. We refer to this signal as $z_{lag,i}^t$ given by:

$$z_{lag,i}^{t+1} = \beta_{lag} z_{lag,i}^t + (1 - \beta_{lag}) z_{avg,i}^t. \quad (14)$$

The absolute difference between the two aforementioned signals determines the information gain of the average signal as:

$$z_{ref,i}^t = |z_{lag,i}^t - z_{avg,i}^t|. \quad (15)$$

As soon as the reference signal is below a certain threshold δ_{prec} the agent increments a time counter and once the counter is above a threshold δ_{mem} it decides to switch to exploitation; i.e. if the reference (gain) signal settles for sufficiently long time close to zero, the agent switches to the exploitation phase. The agent reduces the time counter when the reference signal is above the precision threshold. It is noteworthy that the switching from exploration to exploitation is irreversible, i.e., it can happen only once per agent.

III. METRICS AND SETUP

We assess the performance and speed of collective estimation via different metrics to evaluate the performance of the collective estimation process.

A. Accuracy Metrics

To define what is the exact meaning of accuracy of a CDM, we consider the collective decision as multiple trials to estimate a reference value, which is not easily measurable nor observable. We need to assess how accurately these multiple estimations can capture the reference value. There are three terms that can represent the uncertainty of a CDM [10]. We measure the lack of (collective) accuracy by defining metrics that quantify the accuracy error. The following definitions assume: the memory signal represents the agent's estimate $\hat{z}_i = z_{m,i}$; the collective mean is the arithmetic mean of agents' estimates $\bar{z}_{\text{col}} = \sum_{i=1}^N \hat{z}_i / N$; and the true value is the average intensity in the environment $z_{\text{gt}} = \bar{z}_{\text{env}}$.

1) *Trueness Error*: This metric indicates the overall bias of the collective estimate with respect to the true value. To calculate the trueness error E_T , we need to know what is the true value, which in real-world applications is hard to obtain. Assuming the availability of such information, the trueness error quantifies the distance of the collective mean to the true value. We define:

$$E_T = (\bar{z}_{\text{col}} - z_{\text{gt}})^2. \quad (16)$$

2) *Precision Error*: The closeness of agents' estimations (to each other) is described by precision error. It indicates how diverse are the estimations (or opinions) within the collective, thus the definition of variance quantifies the precision in CDM. We define:

$$E_P = \frac{1}{N} \sum_{i=1}^N (\hat{z}_i - \bar{z}_{\text{col}})^2. \quad (17)$$

3) *Accuracy Error*: The combination of the previous two definitions, that is, the closeness of agents' estimates to the truth, is called accuracy or total error. Other terms such as "effectiveness" of a predictor [33], or "generalization error" [34] have been used for the same concept of accuracy by researchers of different fields. The accuracy error E_A is formulated as:

$$E_A = \frac{1}{N} \sum_{i=1}^N (\hat{z}_i - z_{\text{gt}})^2. \quad (18)$$

4) *Relation between accuracy metrics*: Discriminating between these three definitions in previous works on CDM is rarely discussed. In estimation systems without a bias, precision is exactly equivalent to accuracy. However, the assumption of an unbiased system is questionable, especially given the overall estimation uncertainties. The origin of these errors are also different. The trueness error quantifies the overall bias due to systematic estimation errors. The precision error emerges from all sorts of effects that lead to variability in individual estimates. In the context of collective estimation, it can be regarded as the diversity of estimations. Finally, the accuracy error is the total error representing the estimation uncertainty that entails both errors. It can be shown that by decomposing the total error into bias and variance, the relation between accuracy, trueness, and precision error reads [34]:

$$E_A = E_P + E_T. \quad (19)$$

This equation is fundamental to understand how diversity of estimates is the key parameter controlling accuracy and trueness. In the exploration phase, increasing the diversity of estimates diminishes the trueness error by increasing the precision error and decreasing the accuracy error ($E_T = E_A - E_P$). Whereas in the exploitation phase, decreasing the diversity while keeping the trueness error low reduces the accuracy error. Note that in some figures of section IV, we used the square root of the errors to improve the visualization of our results.

B. Speed Metric

This metric quantifying the speed of the decision process is defined via the duration of the exploration phase. In the pre-defined switching time setup, the decision time is the same as the switching time. While, in the adaptive method, this metric shows how long it takes on average for agents to satisfy the switching condition.

C. Experiment Setup

The experiment setups are designed in a way that the collective needs to explore in order to gather enough information from the environment, in other words, the initial trueness error of the collective is non-zero. Initially, the agents are distributed randomly in a small patch. Therefore, to reduce the trueness error, agents need to explore and

TABLE I: Experiment Setup Parameters

Parameter	Definition	Value
N	Collective Size	100
A	Arena size	[1.4×1.4]
A ₀	Initial placement patch	[0.7×0.7]
d	Communication range	0.30
r _ψ	Change rate of random orientation	0.1
r _λ	Random walk coefficient	0.25
λ	Step size	0.002
σ	Intensity noise coefficient	0.025
α	Weighting factor on memory	0.99
β	Decaying factor for gradient descent	0.99
β _{lag}	Decaying factor for lag signal	0.9
δ _{mem}	Time counter threshold	100

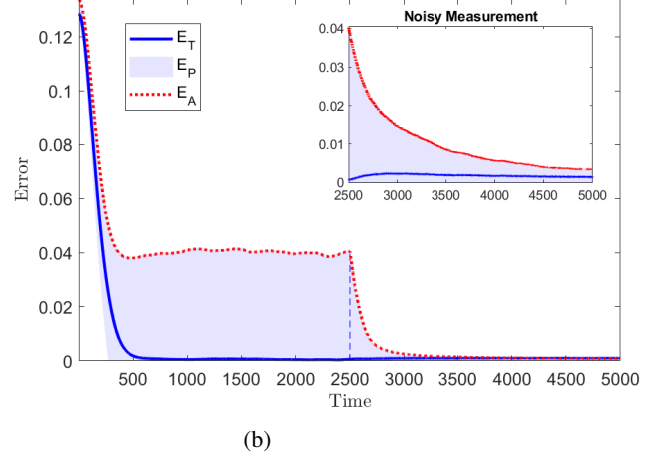
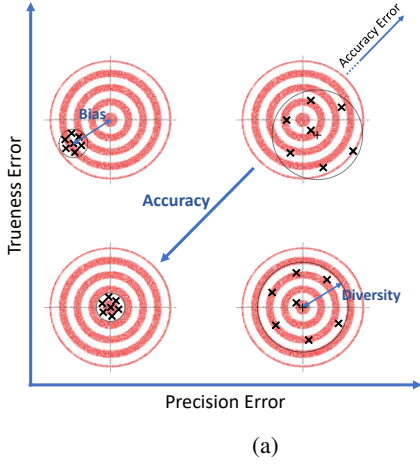


Fig. 3: a) Relationship between trueness, precision, and accuracy error. The collective estimation can be considered as multiple uncertain estimation of a true value, b) Time series of trueness error E_T (solid blue line) showing the bias of the collective estimation. The shade indicates the deviation over the individuals, i.e., the width of the upper half of the shade equals to the precision error E_P . The total accuracy error E_A is the sum of trueness and precision errors (red dotted line). Combining the previous metrics of trueness and precision, this definition of accuracy is informative enough to declare the evaluation of the CDM performance. The dashed line denotes the switching time. All the reported variables here are the average of 40 Monte-Carlo repetitions.

expand the diversity of estimations about the intensity of the environmental feature. The general setup of the numerical experiments is summarized in Table I. The limited requirement for the agents supports the possibility of the method to be implemented on real swarm robots with limited abilities, such as Kilobot [35]. We use two different types of environmental intensity fields or “benchmark functions”, a uni-modal and a multi-modal one, in order to assess the capability of the method for contour finding in various environments (Figure 1). The snapshots of two sample distribution functions are depicted in Figure 1. Note that in the next sections, we discuss about the results only for the uni-modal distribution. Using equations 8, 11, and 12, we simulated the agents as particles moving in a two-dimensional bounded environment.

IV. RESULTS AND DISCUSSION

A. Pre-defined Switching Time

Figure 3-b shows for the non-adaptive method in the exploration and exploitation task. For small times, prior to the switching time $t = 2500$ indicated by the vertical dashed line, we observe two phenomena: 1) The approach of the mean value perceived by the collective to the true environmental average, and 2) an increasing diversity of the collective information (shown by the shade.) These results indicate that the driving force behind the reduction of the trueness error is the increase in estimate diversity. Generally speaking, the more evidence the agents gather from the environment and contribute to the collective knowledge, the less biased is the resulting collective estimation. During the first phase, increasing the diversity (decreasing the precision cf. Figure 3-a) increases the trueness of the estimation. Similarly, the accuracy of the collective estimation also increases

(reducing E_A) during the exploration. While the trueness error converges asymptotically to zero, the accuracy error stops to decrease and saturates at a finite value. Considering the bias-variance decomposition of the total error (Eq. 19), the remained accuracy error is due to the precision error. For the collective to reduce the precision error, the agents start to exploit the information by communicating and aggregating the information.

In contrast to the exploration phase, where agents contribute to enriching the collective wisdom by providing diverse pieces of information, in the exploitation phase, it is the collective wisdom that contributes to the estimation of individuals. To elaborate more, consider the estimation accuracy of a random agent at the end of the exploration phase. It can be any random variable that is measurable in the arena. However, at the end of the exploitation phase, each individual’s estimation is roughly as accurate as that of any other individual in the collective, which is also as accurate as the estimation of the collective as a whole. With that in mind, the second phase is where the agents are able to exploit the potential information contained within the collective, which, thanks to the exploration phase, is now less biased. Thus, the exchange of the information increases the collective’s precision and, as a result, increases the accuracy of the collective estimation. The trueness error, however, remains unchanged. Finally, it can be deduced that at the end of the task, each agent can estimate accurately, as if it has access to the full collective information. It means that if the agents are meant to decide to do an operation based on the average intensity of a feature in the arena, they are able to make the decision accurately, and in a decentralized manner.

B. Switching-time vs Accuracy

Since having a priori knowledge about the switching time is not practical or easy to obtain, it is essential to assess how varying the switching time (t_{sw}) affects estimation accuracy. We evaluated the accuracy error of the collective estimation for different switching times, while keeping the overall duration of the experiments fixed ($t_f = 5000$). As shown in Figure 4, increasing the switching time from zero, i.e. extending the time spent in the exploration phase, results initially in a drastic decrease in accuracy error, which is a consequence of the reduced bias of the collective estimation. We associate this to the speed-vs-accuracy tradeoff in decision making, stating that the collective has to sacrifice the accuracy if individuals decide to switch early. Negative effects of late switching become apparent when the time budget for the whole task is limited. In this case, spending too much time on the exploration restricts the time for the aggregation both in information and spatial domains. This again decreases the accuracy of the CDM (for $t_{sw} > 2500$), due to the higher values of precision error.

To clarify how this SAT affects the trueness, precision and accuracy errors of the CDM we put the *inset* graph showing the trueness error for the same experiments as the main figure. Considering that the total error (E_A) is the combination of the precision and trueness error (Eq. 19), we can deduce how the switching time can change the contribution of the precision and trueness error on the total accuracy error. These results reveal an important phenomenon in the SAT paradigm of CDM. In contrast to individual decision-making, where decision accuracy monotonically increases by time, in CDM that is not the case. Because, by changing the balance between exploration and exploitation, the share of the two antagonistic/competitive components of the total accuracy (trueness and precision) changes.

Moreover, we investigated three different arena sizes and our results show that the arena size plays a significant role in the SAT, especially with pre-defined switching time. This study shows that for the collective to obtain a certain amount of accuracy in larger arenas, it has to spend more time exploring. The optimal switching time for the collective increases with increasing arena size as expected.

C. Adaptive Switching Time

Given the difficulty to identify an optimal switching time in unknown environments, we discuss here the performance of the system with the individual-level adaptive approach as introduced above. The results on the collective estimation performance with the adaptive switching mechanism for different arena sizes is illustrated in Figure 5. As for the fixed switching time case, the initial agent positions are randomly distributed in the same small patch in the differently sized arenas. As a result, the initial estimation error is different for each arena size; which is the reason why in larger arenas the collective needs to spend more time in the exploration phase for the same reduction of the bias error. The *inset* of that figure shows that in larger arenas, on average, it takes more time for the collective to switch to the exploitation phase.




Fig. 4: Accuracy error versus switching time: Each curve shows the steady-state accuracy error, averaged over 40 Monte Carlo simulations, for various switching time for a given arena size. The solid and dotted lines are for simulations where agents have perfect ($\sigma = 0$) and noisy measurement ($\sigma = 0.025$), respectively. The black line: small arena [1×1], the blue line: medium arena [1.4×1.4], the red line: large arena [1.73×1.73].

To study the underlying speed-vs-accuracy tradeoff, we varied the adaptive method parameter δ_{prec} and ran 40 independent Monte-Carlo simulations. The average decision time and accuracy error were obtained for each parameter value. In Figure 6, each symbol indicates the performance of the collective with respect to speed and accuracy for a certain value of the adaptive parameter δ_{prec} , with the increasing lightness of the colored symbols representing lower values of δ_{prec} . To explore the effect of limited time budgets, we compare results for different termination (final) times t_f of the simulation. The green symbols depict the results for $t_f = 5000$, while for the purple symbols the time budget for the complete scenario is increased by a factor of 10 by setting $t_f = 50000$. Increasing δ_{prec} makes the agents to switch earlier to the exploitation phase, which for small δ_{prec} ($t_{sw} \lesssim 2000$) results in a lower accuracy error. It is qualitatively the same conclusion as in the fixed switching case (Figure 4): Investing insufficient amount of time into exploration leads to a larger collective bias, as measured by the trueness error (see inset Fig. 4).

Comparing the performance curves (in Figure 6) for different t_f for small values of δ_{prec} (i.e., large average switching times) shows a clear impact of limited time budget on the SAT for the adaptive method. For low total simulation time, we observe a U-shape curve with a strong decrease in accuracy (increase in E_A) for small δ_{prec} , corresponding to long durations of the exploration phase. The trueness error in this parameter regime is low independent of the total simulation time (*inset* Figure 6). The source of this overall increase in accuracy error for limited time budget with decreasing δ_{prec} ($t_{sw} \gtrsim 2000$), can be traced back to an increase in precision error due to the limited and decreasing time available for the exploitation phase. As in Figure 4,

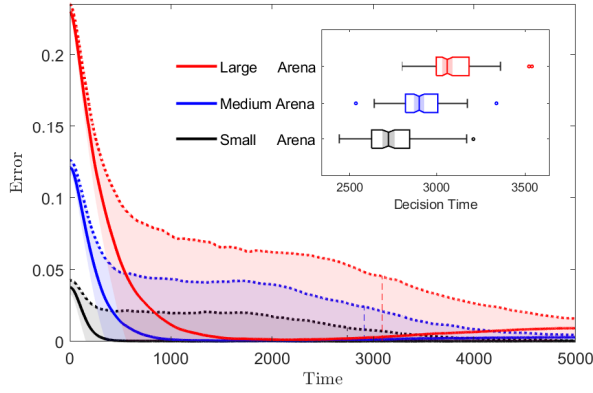


Fig. 5: The performance of the adaptive method with fixed parameters in different arena sizes for $\delta_{\text{prec}} = 10^{-6}$. The black, blue, and red lines correspond to small, medium, large arena sizes, respectively. The shades show the deviation of agents from the mean value, i.e., the precision error E_P . The dotted lines are the accuracy error E_A , and the dashed lines indicate the mean decision time of the collective. To make the decision times more clear, the box plot of the inset is provided.

the elevated tail of the speed-vs-accuracy graph highlights the importance of aggregation time for collective decision-making.

V. CONCLUSION

We have studied a collective estimation problem in a continuous spatial environment with moving agents. Agents have the ability to measure a noisy signal from the environment and communicate with their local neighbors by exchanging their opinion about the intensity of a specific feature in the environment. We propose a solution focused on exploration-exploitation, where the switching time was considered as a control parameter. The aim of the exploration phase is to disperse the collective in the arena and collect diverse measurements, while there is no exchange of information during this phase. Avoiding agents to share their information during exploration helps the collective by preventing the spread of misestimation throughout the network. In contrast, the collective aggregates during the exploitation phase, both in the information and spatial domain. By decomposing the total accuracy error into bias and variance (Eq. 19), we discussed how diversity is the controlling factor, which itself is modulated by either the random independent movements or social interaction and collective motion.

We showed that the switching time plays an important role in changing the balance between exploration and exploitation, which directly affects the tradeoff between speed and accuracy of the collective estimation. We highlighted the importance of considering systematic bias and random variance components of total error, especially when it comes to understanding the performance of the collective in the context of a SAT. We also observed the SAT on different levels. During

the exploration and exploitation phases we observed decreasing trueness and precision errors over time, respectively, and we varied the switching time, which changes the exploration-exploitation balance. This explains that in CDM, the SAT pattern arises not only because of gathering more information during the exploration phase [24], but also because consensus is achieved by time. We stressed that the SAT paradigm in CDM may exhibits higher complexity than in individual DM. The difference is more critical when it comes to a limited time budget for the whole scenario. The results proved that making a decision slower (greater switching time) results in a reduced trueness error of the collective estimation, while increases precision error. By contrasting the precision and trueness errors, we showed that aggregation of information is a time-consuming process, and consequently, the precision error, i.e. the error in spatial positioning, cannot be regulated instantaneously, but is subject to spatio-temporal constraints set by the individual movements.

Furthermore, the ability of agents to move in the spatial domain, and the definition of the objective function, allows the collective to move collectively and to capture the iso-contours of the intensity distribution. There is a variety of potential applications for such a fully distributed collective behavior with spillage capturing and source localization [36], being two examples. The proposed mechanisms, for the interaction and motion, have the ability to deal with a dynamic environment. We, however, limited our focus to a static environment.

REFERENCES

- [1] H. Hamann, *Swarm robotics: A formal approach*. Springer, 2018.
- [2] F. Galton, “Vox populi,” 1907.
- [3] I. Yaniv and M. Milyavsky, “Using advice from multiple sources to revise and improve judgments,” *Organizational Behavior and Human Decision Processes*, vol. 103, no. 1, pp. 104–120, 2007.
- [4] J. Surowiecki, *The wisdom of crowds*. Anchor, 2005.




Fig. 6: Accuracy error versus switching time for various δ_{prec} parameter of the adaptive switching method. Each point is the average of 40 independent Monte-Carlo simulations for each parameter value. The total duration of the simulation for the green points is 5000 time-steps, whereas that of the purple points is 50000.

- [5] L.-A. Giraldeau, T. J. Valone, and J. J. Templeton, "Potential disadvantages of using socially acquired information," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 357, no. 1427, pp. 1559–1566, 2002.
- [6] A. Novaes Tump, M. Wolf, J. Krause, and R. H. Kurvers, "Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information," *Journal of the Royal Society Interface*, vol. 15, no. 142, p. 20180155, 2018.
- [7] C. Winklmayr, A. B. Kao, J. B. Bak-Coleman, and P. Romanczuk, "The wisdom of stalemates: consensus and clustering as filtering mechanisms for improving collective accuracy," *Proceedings of the Royal Society B*, vol. 287, no. 1938, p. 20201802, 2020.
- [8] L. Hong and S. E. Page, "Groups of diverse problem solvers can outperform groups of high-ability problem solvers," *Proceedings of the National Academy of Sciences*, vol. 101, no. 46, pp. 16385–16389, 2004.
- [9] S. E. Page, *The difference: How the power of diversity creates better groups, firms, schools, and societies-new edition*. Princeton University Press, 2008.
- [10] J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing, "How social influence can undermine the wisdom of crowd effect," *Proceedings of the national academy of sciences*, vol. 108, no. 22, pp. 9020–9025, 2011.
- [11] J. Becker, D. Brackbill, and D. Centola, "Network dynamics of social influence in the wisdom of crowds," *Proceedings of the national academy of sciences*, vol. 114, no. 26, pp. E5070–E5076, 2017.
- [12] A. Novaes Tump, "How social learning strategies boost or undermine decision making in groups," Ph.D. dissertation, Freien Universität Berlin, 2020.
- [13] W. A. Wickelgren, "Speed-accuracy tradeoff and information processing dynamics," *Acta psychologica*, vol. 41, no. 1, pp. 67–85, 1977.
- [14] L. Chittka, P. Skorupski, and N. E. Raine, "Speed-accuracy tradeoffs in animal decision making," *Trends in ecology & evolution*, vol. 24, no. 7, pp. 400–407, 2009.
- [15] N. R. Franks, A. Dornhaus, J. P. Fitzsimmons, and M. Stevens, "Speed versus accuracy in collective decision making," *Proceedings of the Royal Society of London. Series B: Biological Sciences*, vol. 270, no. 1532, pp. 2457–2463, 2003.
- [16] V. Srivastava and N. E. Leonard, "Collective decision-making in ideal networks: The speed-accuracy tradeoff," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 121–132, 2014.
- [17] R. Bogacz, E.-J. Wagenmakers, B. U. Forstmann, and S. Nieuwenhuis, "The neural basis of the speed-accuracy tradeoff," *Trends in neurosciences*, vol. 33, no. 1, pp. 10–16, 2010.
- [18] M. M. Sosna, C. R. Twomey, J. Bak-Coleman, W. Poel, B. C. Daniels, P. Romanczuk, and I. D. Couzin, "Individual and collective encoding of risk in animal groups," *Proceedings of the National Academy of Sciences*, vol. 116, no. 41, pp. 20556–20561, 2019.
- [19] A. S. Wade, I. W. Ramnarine, and C. C. Ioannou, "The effect of group size on the speed of decision making depends on compromise and predation risk across populations in the guppy *poecilia reticulata*," *Behaviour*, vol. 1, no. aop, pp. 1–20, 2020.
- [20] G. Valentini, H. Hamann, and M. Dorigo, "Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off," in *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 2015, pp. 1305–1314.
- [21] G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo, "Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems," *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 3, pp. 553–580, 2016.
- [22] J. T. Ebert, M. Gauci, F. Mallmann-Trenn, and R. Nagpal, "Bayes bots: collective bayesian decision-making in decentralized robot swarms," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 7186–7192.
- [23] M. S. Talamali, J. A. Marshall, T. Bose, and A. Reina, "Improving collective decision accuracy via time-varying cross-inhibition," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 9652–9659.
- [24] T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, I. D. Couzin, C. S. R. Group *et al.*, "Exploration versus exploitation in space, mind, and society," *Trends in cognitive sciences*, vol. 19, no. 1, pp. 46–54, 2015.
- [25] P. Landgren, V. Srivastava, and N. E. Leonard, "Distributed cooperative decision-making in multiarmed bandits: Frequentist and bayesian algorithms," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 167–172.
- [26] —, "Distributed cooperative decision making in multi-agent multi-armed bandits," *Automatica*, vol. 125, p. 109445, 2021.
- [27] K. Mehlhorn, B. R. Newell, P. M. Todd, M. D. Lee, K. Morgan, V. A. Braithwaite, D. Hausmann, K. Fiedler, and C. Gonzalez, "Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures," *Decision*, vol. 2, no. 3, p. 191, 2015.
- [28] P. Romanczuk, M. Bär, W. Ebeling, B. Lindner, and L. Schimansky-Geier, "Active brownian particles," *The European Physical Journal Special Topics*, vol. 202, no. 1, pp. 1–162, 2012.
- [29] A. Okubo and S. A. Levin, *Diffusion and ecological problems: modern perspectives*. Springer, 2001, vol. 14.
- [30] M. Zellner, A. E. Abbas, D. V. Budescu, and A. Galstyan, "A survey of human judgement and quantitative forecasting methods," *Royal Society Open Science*, vol. 8, no. 2, p. 201187, 2021.
- [31] A. M. Simons, "Many wrongs: the advantage of group navigation," *Trends in ecology & evolution*, vol. 19, no. 9, pp. 453–455, 2004.
- [32] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.
- [33] S. Geman, E. Bienenstock, and R. Doursat, "Neural networks and the bias/variance dilemma," *Neural computation*, vol. 4, no. 1, pp. 1–58, 1992.
- [34] A. Krogh and J. Vedelsby, "Validation, and active learning," *Advances in neural information processing systems 7*, vol. 7, p. 231, 1995.
- [35] M. Rubenstein, C. Ahler, and R. Nagpal, "Kilobot: A low cost scalable robot system for collective behaviors," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3293–3298.
- [36] A. A. R. Newaz, S. Jeong, H. Lee, H. Ryu, N. Y. Chong, and M. T. Mason, "Fast radiation mapping and multiple source localization using topographic contour map and incremental density estimation," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1515–1521.