




# Best-of- $N$ Collective Decisions on a Hierarchy

Fabio Oddi<sup>1</sup>, Andrea Cristofaro<sup>1</sup> , and Vito Trianni<sup>2</sup>  

<sup>1</sup> DIAG, Sapienza University of Rome, Rome, Italy  
`fabio.oddidi@uniroma1.it`, `cristofaro@diag.uniroma1.it`

<sup>2</sup> ISTC, National Research Council, Rome, Italy

**Abstract.** The best-of- $N$  problem in collective decision making is complex especially when the number of available alternatives is larger than a few, and no alternative distinctly shines over the others. Additionally, if the quality of the available alternatives is not a priori known and noisy, errors in the quality estimation may lead to the premature selection of sub-optimal alternatives. A typical speed-accuracy trade-off must be faced, which is hardened by the presence of several alternatives to be analyzed in parallel. In this study, we transform a one-shot best-of- $N$  decision problem in a sequence of simpler decisions between a small number of alternatives, by organizing the decision problem in a hierarchy of choices. To this end, we construct an  $m$ -ary tree where the leaves represent the available alternatives, and high-level nodes group the low-level ones to present a low-dimension decision problem. Results from multi-agent simulations in both a fully-connected topology and in a spatial decision problem demonstrate that the sequential collective decisions can be parameterized to maximize speed and accuracy against different decision problems. A further improvement relies on an adaptive approach that automatically tunes the system parameters.

## 1 Introduction

Collective decisions are key in many natural and artificial systems, from bacteria to social insects, from organizations to robot swarms [11, 18, 22, 24]. A collective decision requires that a group of agents agrees on a common solution to a given problem. Generally speaking, a best-of- $N$  collective decision problem presents multiple alternative solutions, each characterized by a quality value representing associated benefits and costs, and the best possible one must be chosen by a (qualified) majority or through full consensus. Intuitively, the larger is the number  $N$  of alternatives, the harder is the problem of selecting the best one. This is because a qualified majority is more difficult to form if individual choices can spread across many competing alternatives, and because correctly evaluating and comparing multiple alternatives becomes extremely cumbersome and can lead to spreading of individual biases. Indeed, previous studies on collective decision making revealed that increasing the number of alternatives negatively affects the

ability to make accurate decisions. First, the parameter space in which a decision deadlock is possible is larger when  $N$  increases, suggesting that the system may remain unable to collectively choose any alternative for a long time, and that sub-optimal alternatives could be selected [16]. Second, the time taken to select the best alternative increases exponentially with the number of available alternatives, in adherence with the Hick-Hyman’s law of psychophysics [15].

In the face of the complexity of deciding among multiple alternatives, there is evidence that the path towards decision does not directly lead to a crisp choice of the best alternative [7]. Instead, it is likely that one or a group of alternatives is discarded quickly to reduce the problem dimensionality. This is the case for choices among spatially distributed alternatives in both individuals and collectives, where the feedback between body movements and decision making allows to select out alternatives when their reachability decreases [23]. Similarly, optimal decision heuristics in multi-alternative choice problems allow to focus attention on a limited number of alternatives, often as small as two, reducing the cognitive load for a comprehensive evaluation [6]. When evidence must be gathered about the available alternatives, a speed-accuracy trade-off arises, where higher accuracy can be achieved by collecting additional evidence, which slows down decision making [5]. Selective attention to different alternatives can be modulated according to value, which implies that low-quality alternatives are sampled less often [9]. On the other hand, adaptive sampling of alternatives postulates that increased sampling is dedicated to alternatives when quality uncertainty is higher [8]. Both these mechanisms suggest that sequential sampling and evidence accumulation in multi-alternative decision problems can lead to focusing on just the most interesting options, putting aside less valuable ones.

Correlations between alternatives (e.g., spatial location) may require optimal planning of the sampling effort. This problem has been addressed in robotics mapping, where a multi-resolution adaptive sampling can optimize evidence accumulation [10]. On the basis of the gathered evidence, (collective) decision making can lead to optimal deployment of robots in the area where working is most valuable [1, 2]. In these studies, an  $m$ -ary tree is introduced to represent the spatial correlation between alternatives (i.e., areas where some work is required). This representation also drives the resolution of the evidence accumulation, enabling to observe multiple alternatives at the same time or to focus on subgroups.

Inspired by the above considerations, we study whether collective decision making can benefit from a hierarchical organization of the  $N$  available alternatives, which can direct the evidence accumulation and decision process. We choose a hierarchical organization as the most straightforward way of aggregating alternatives, by maintaining possible correlations (e.g., closeness in space) and recursively creating macro-alternatives that group together some of the available options. Our working hypothesis is that a complex decision process can be simplified by reducing it to a sequence of decisions between a smaller number of groups of alternatives. Evidence accumulation is performed at the group level, hence with a reduced resolution (i.e., aggregating randomly sampled val-

ues within a group). Decisions are collectively taken following a decentralized algorithm inspired by house-hunting honey bees [17], which can be controlled by a single parameter describing the ratio of interactions with respect to individual decision. We complete the algorithm with quorum sensing, which enables to efficiently serialize the decision making process over the defined hierarchy [14, 21]. We show that the hierarchical approach allows to solve different decision problems with the same parameterization, while the non-hierarchical approach does not present a single parameterization that both maximizes accuracy and is capable of breaking deadlocks among equivalent alternatives. These results are robust to noisy evidence accumulation—owing to the possibility of agents to share their local knowledge and collectively minimize the estimation error—and to local quorum estimation. We also introduce an adaptive parameterization of the collective decision-making algorithm based on the residual uncertainty about the quality of alternatives. Thanks to this mechanism, it is possible to improve the speed-accuracy trade-off for all tested hierarchies. Finally, we demonstrate that the proposed approach can be adapted well to spatial best-of- $N$  decisions.

## 2 Experimental Setup

### 2.1 Problem Description

In best-of- $N$  decisions,  $N$  different alternatives are present, each characterized by a quality value  $v_i$ ,  $i \in [1, N]$ . A group of  $M$  agents must collectively choose the alternative with the highest quality, or one of the equal-best alternatives. Without loss of generality, we assume that one alternative has maximum value (here,  $v_M = 10$ ), while all other alternatives have the same, small value ( $v_m = \kappa v_M$ ,  $\kappa \in [0, 1]$ ). We consider that a decision has been made by the group of agents when a quorum  $Q \in (0.5, 1]$  is reached for one of the alternatives, that is, a minimum fraction  $Q$  of agents has selected the same alternative.

In our hierarchical problem formulation, the alternatives are organized in a  $m$ -ary tree, that is, a tree where each parent node has at most  $m$  children. Each alternative corresponds to a leaf node in the tree, while non-leaf nodes represent groups of alternatives. Given  $N$  alternatives, the minimum depth of an  $m$ -ary tree is  $D = \lceil \log_m N \rceil$ . Hence, a binary tree with depth  $D = 5$  contains at most  $N = 32$  alternatives. Conversely, when  $m = N$ , the depth of the tree is  $D = 1$  meaning that there is no hierarchical organization of the alternatives, which are all children of the root node. A node at level  $d \in \{0, \dots, D\}$  of the tree entails a decision between at most  $m$  alternatives. We consider here only perfect  $m$ -ary trees, where each non-leaf node has exactly  $m$  child nodes and all leaf nodes are at the same depth. A tree node is labeled by  $n_j^d$ , where  $d$  is the depth and  $j \in 1, \dots, m^d$  is a univocal index. The sub-tree defined by  $n_j^d$  is referred to as  $\mathcal{S}_j^d$ . The root node ( $n_1^0$ ) is the starting point for collective decision making.

### 2.2 Collective Decision Process

The collective decision process extends a design pattern for decentralized decision making [17] with quorum sensing to move at different tree depth. We start by

describing the ideal case in which agents have knowledge about the number and quality of the alternatives, as well as about the number of agents that are in their same sub-tree. Later, we relax these assumptions.

**Agents with Perfect Knowledge.** In the hierarchical collective decision model, at any time  $t$  an agent  $i$  is characterized by a tuple  $a_i = \langle s_i, c_i \rangle$ , where  $s_i$  is the node in which the agent resides, and  $c_i$  is a desired destination node. If  $c_i = s_i$ , the agent is said to be “uncommitted at  $s_i$ ”, that is, the agent has not selected a desired destination. Otherwise,  $c_i$  can be any children of  $s_i$  and the agent is said to be “committed to  $c_i$ ”. At start, all agents are initialized with the tuple  $\langle n_1^0, n_1^0 \rangle$ , that is, they are uncommitted at the root node. We denote with  $\mathcal{P}_j^d$  the sub-population of agents residing at any node in the sub-tree starting from  $n_j^d$ . Note that  $\mathcal{P}_1^0$  is the entire agent population, with  $|\mathcal{P}_1^0| = M$ .

Agents change their residence node only if a quorum of committed agents has been reached for the child node they are committed to. Specifically, an agent  $i$  residing in node  $s_i = n_j^d$  and committed to the child node  $c_i = n_l^{d+1}$  moves its residence to the latter if  $|\mathcal{P}_l^{d+1}| \geq QM$ . Once such a quorum is reached, all committed agents change their residence node, ensuring that a large fraction of the population moves down the hierarchical structure towards the leaf nodes. This also means that agents never need to move back to the parent node (but see below when imperfect quorum sensing is implemented). The process ends if a quorum is reached for one of the leaf nodes representing one of the available options. In other words, agents remain in their residing node unless a quorum is reached, and then move to the selected child node. Given that a sufficiently large quorum is the result of a collective decision, the agent population is expected to perform a sequence of  $D$  decisions leading to the selection of one leaf node.

The commitment state of an agent changes according to stochastic processes inspired by the decision process of house-hunting honeybees [17]. In particular, we exploit the parameterization extensively studied in [15, 16]. To this end, the quality of a node must be considered, which must be a function of the group of alternatives that a node represents. Here, we consider that the quality of a node  $v(n_j^d)$  is recursively computed as the maximum quality among the child nodes:

$$v(n_j^d) = \max_{n \in \mathcal{C}(n_j^d)} v(n), \quad (1)$$

where  $\mathcal{C}(n_j^d)$  is the set of child nodes of  $n_j^d$ . This allows to propagate up in the hierarchy the best value of the underlying alternatives represented in the leaf nodes, without loosing in resolution. See Sect. 4 for a discussion.

At time  $t$ , an agent  $i$  uncommitted at  $n_j^d$  ( $a_i = \langle n_j^d, n_j^d \rangle$ ) can spontaneously become committed to a randomly selected child node  $n_l^{d+1}$  with probability:

$$P_\gamma = k \frac{v(n_l^{d+1})}{v_M}, \quad (2)$$

where  $k$  is a tunable parameter chosen to scale the probability. At the same time, the agent  $i$  may get recruited by another agent  $b \in \mathcal{P}_j^d$ , with  $a_b = \langle s_b, c_b \rangle$ . This is

implemented by choosing a random agent from  $\mathcal{P}_j^d$  and computing the following recruitment probability:

$$P_\rho = \begin{cases} h \frac{v(c_b)}{v_M} & c_b \neq n_j^d \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $h$  is a scaling factor. Note that commitment and recruitment are evaluated in parallel, requiring that  $P_\gamma + P_\rho \leq 1$ . Hence, we impose that  $h + k = 1$ .

When an agent  $i$  is committed to  $n_l^{d+1}$ , it can spontaneously become uncommitted at  $n_j^d$  with an abandonment probability inversely proportional to quality:

$$P_\alpha = k \frac{1}{1 + v(n_l^{d+1})}, \quad (4)$$

where  $k$  is the same scaling factor as in (2). Also, the agent  $i$  may get inhibited by another agent  $b \in \mathcal{P}_j^d$ , with  $a_b = \langle s_b, c_b \rangle$ . This is implemented by choosing a random agent from  $\mathcal{P}_j^d$  and computing the inhibition probability:

$$P_\sigma = \begin{cases} h \frac{v(c_b)}{v_M} & c_b \neq n_j^d \wedge c_b \neq n_l^{d+1} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where  $h$  is the same scaling factor as in (3). Also in this case, we must enforce that  $P_\alpha + P_\sigma \leq 1$ , which is still the case if we ensure that  $h + k = 1$ .

Note that, with a flat hierarchy ( $D = 1$ ), the collective decision process reduces to the standard best-of- $N$  decisions previously studied [15, 16]. In analogy with previous work, we introduce a single control parameter  $r = \frac{h}{k}$  to tune the relative strength of stochastic processes based on interactions with other agents (i.e., recruitment and inhibition) with respect to spontaneous stochastic processes (i.e., commitment and abandonment).

**Estimation of the Alternative Quality.** Moving beyond the ideal case presented above, we introduce here the more realistic case in which agents are aware about the maximum number  $N$  of alternatives and their hierarchical organization, but are not aware of their quality, which is perceived with noise. At every decision step, an agent  $i$  makes an observation of a node  $n$  chosen according to its state: if the agent is uncommitted at  $n_j^d$ , it observes the child node selected for the computation of the commitment probability (2). Otherwise, the agent observes the desired destination  $c_i$ . If the observed node  $n$  is a non-leaf node, then a random leaf in the sub-tree of  $n$  is chosen for observation, and all parent nodes are updated according to (1). Upon observation of a leaf node  $n_j^D$ , agent  $i$  updates its quality estimate  $\tilde{v}_i$  according to a moving average:

$$\tilde{v}_i(n_j^D) \leftarrow \begin{cases} v(n_j^D) + N(0, \sigma_j^D) & t_j^D = 0 \\ \lambda \tilde{v}_i(n_j^D) + (1 - \lambda)(v(n_j^D) + N(0, \sigma_j^D)) & \text{otherwise} \end{cases}, \quad (6)$$

where  $t_j^D$  counts the number of observations previously performed,  $\lambda$  represents the smoothing factor and  $\sigma_j^D$  represents the variance of a Gaussian noise. This

simple exponential filter allows to rapidly converge on a stable estimation of the quality of the alternatives, and also allows to adapt to changing qualities—a possibility not explored in this study. Additionally, by limiting the observations to the sub-tree where an agent resides, the estimation is focused on the relevant alternatives, avoiding to waste time for those that have been discarded earlier.

To take advantage of the collective sensing abilities, agents exchange their current estimates upon interaction: when agent  $i$  interacts with agent  $b$ , it receives from  $b$  the quality estimates of all the  $N$  alternatives. These estimates are treated in the same way as independent observations to update the individual estimates with (6). Overall, this leads to a fast convergence towards the average.

**Quorum Sensing.** If agents cannot reliably count how many agents are committed to a given node, an estimate can be obtained by keeping memory of the last interactions. At every time step  $t$ , an agent  $i$  that interacts with another agent  $b$  records in a list  $\mathcal{L}$  a tuple  $\langle b, c_b, t_b = t \rangle$ . In case an element is present in  $\mathcal{L}$  with the same id  $b$ , the corresponding timestamp  $t_b$  is updated. Finally, old elements are purged from  $\mathcal{L}$  when  $t - t_b > T_M$ , where  $T_M$  is a maximum period for retaining past interactions. As agents have only one interaction per time step, it follows that  $|\mathcal{L}| \leq T_M$ . Quorum sensing is implemented looking at the information in  $\mathcal{L}$ . An agent  $i$  committed to  $n_i^{d+1}$  counts the number  $L_i^{d+1}$  of elements in the list where  $c_b \in \mathcal{S}_i^{d+1}$ . If  $L_i^{d+1}$  is larger than a threshold  $L_M$ , the agent has recently interacted with a sample of the population with the same commitment. Hence, the agent considers the quorum reached and changes its residing node to  $s_i = n_i^{d+1}$ . This estimation is however prone to errors, because of small samples or because old information does not represent any more the current population. Hence, a recovery mechanism allows robots to move to the parent node. Given that  $\mathcal{L}$  shrinks if an agent resides in a node in which the population is small (e.g., when a real quorum was not actually reached), an agent uncommitted at  $n_j^d$  changes its residence to the parent node if  $|\mathcal{L}| < L_m$ , where  $L_m < L_M \leq T_M$ .

**Adaptive Parameter Selection.** The parameter  $r = \frac{h}{k}$  can determine if and how fast the group converges to a shared choice [16]. High values correspond to fast but possibly inaccurate decisions, as social information is given more importance than individual quality estimation. To improve the decision making process, social feedback should increase when the uncertainty about the available options decreases. Agents measure uncertainty by learning two independent Gaussian models for each alternative—i.e., updating one model every second observation. Then, the Hellinger distance  $H_d \in [0, 1]$  between the two is computed and associated to the corresponding leaf node. Non-leaf nodes receive the maximum distance of their children. High  $H_d$  corresponds to insufficient sampling, hence high uncertainty. An agent  $i$  committed to  $n_i^{d+1}$  uses the  $H_d(n_i^{d+1})$  to compute a value  $r_i = g(1 - H_d)$ , where  $g$  is a gain to adjust the range. During the decision process, an agent  $i$  interacting with agent  $b$  receives from the latter

the value  $r_b$ . This allows to tune the process strength after the uncertainty of the interacting agent, where low uncertainty corresponds to stronger interactions.

**Decision Making on a Spatial Hierarchy.** We also consider the spatial case in which agents have to identify the most valuable area within a given region. The region is divided in  $N$  areas, and a hierarchical organization is imposed recursively partitioning the region in smaller areas. The root node of the hierarchical structure represents the whole region, which is partitioned in  $m$  sub-regions, and each is recursively divided further until  $N$  areas are obtained. Similarly to the non-spatial setup, only one area has maximum quality  $v_M$ , while all other areas have the same quality  $v_m = \kappa v_M$ . We consider here the case of binary and quad-trees, usually employed to represent a 2D space. Agents move according to a random waypoint model [4]: an agent  $i$  selects a random position within the region corresponding to the node  $c_i$  within the hierarchy, and moves there at constant speed. To focus on the effects of spatiality, we ignore collisions, as if the agent body size were negligible with respect to the dimensions covered (e.g., in case of drones monitoring a large field [2]). In the future, this assumption can be relaxed via efficient velocity-obstacle collision avoidance [3].

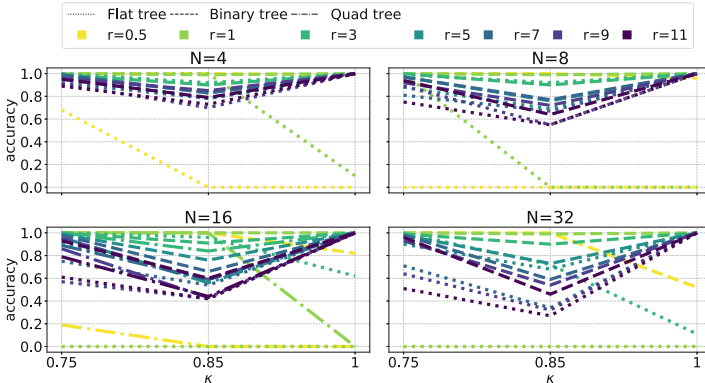
The decision process is implemented through space as follows. An agent  $i$  randomly selects a position in  $c_i$  and moves there. Once the random position is reached, it broadcasts information about its ID  $i$ , its current state  $a_i$ , the current quality estimates for the  $N$  alternatives, and a timestamp  $t_i$ . Broadcast messages can be received within a limited communication radius  $R_c$ . Upon reception of a new message, an agent re-broadcasts it, and stores the information in a list  $\mathcal{N}$  of detected neighbors, overwriting any older message from the same sender. Quality estimates are updated according to the information received from neighbors. Messages older than  $T_M$  are purged from  $\mathcal{N}$ . When the agent reaches the desired position, a noisy observation is made at the agent location, and the tree structure is updated accordingly. Then, the decision process takes place. If the agent is uncommitted at  $n_j^d$ , it computes the commitment probability for the child node  $n_l^{d+1}$  corresponding to the current agent position. A random neighbor is selected from  $\mathcal{N}$  and the recruitment probability is computed. Similarly, if the agent is committed to  $n_l^{d+1}$ , the abandonment probability is computed with the latest quality estimate, and the inhibition probability from the randomly selected neighbor from  $\mathcal{N}$ . In any case, the information from the selected neighbor is used to update the list  $\mathcal{L}$  for quorum sensing, similarly which is implemented as in the non-spatial case. Note that quorum is not computed on the list  $\mathcal{N}$  because to avoid overestimating the opinion of the local population, and to keep the decision process aligned with the non-spatial case. Note also that in the spatial case, decisions are taken only when a new observation is made at the randomly selected position, contrary to the non-spatial case in which decisions occur at any time step. Hence, the spatial case evolves slowly, but an equivalence can be made looking at the average number of decisions made within the population.

### 3 Results

We measure accuracy as the percentage of runs (out of 100) in which the best alternative—or one of the equal-best when  $\kappa = 1$ —was correctly identified by at least  $QM$  agents. As a proxy for decision speed, we measure the convergence time taken by the agents to reach a collective decision. To this end, we employ the Kaplan-Meier estimator [12] to compute the empirical cumulative distribution of convergence times, censoring the runs that do not converge within the maximum allotted time. We fit a Weibull distribution and use the fitted function to compute average and standard deviation of the convergence times.


First of all, we consider the non-spatial case in which agents have perfect knowledge. We consider a relatively easy decision problem with  $\kappa = 0.75$ , a more difficult one with  $\kappa = 0.85$ , and a symmetry breaking problem with  $\kappa = 1$ . The parameter  $r = \frac{h}{k}$  is used as control parameter to tune convergence speed. Figure 1 shows how the accuracy varies across problems for different configurations were flat, binary and quad-trees are employed (the latter only for  $N = 16$ ). For any value of  $N$  we tested, the binary tree provides the highest accuracy when  $r = 1$ . A lower value ( $r = 0.5$ ) is equally good unless  $N \geq 16$  and  $\kappa = 1$ , where several runs do not terminate within the allotted time  $\mathcal{T}$ . Indeed, when  $r$  is small, the collective decision process cannot take full advantage of the positive and negative feedback loops, making symmetry breaking very difficult. Conversely, higher values of  $r$  may excessively rely on social information, which may lead to a loss in accuracy when the decision problem is difficult ( $\kappa = 0.85$ ).

Non-hierarchical structures ( $m = N$ ) struggle to consistently provide good results across all problem configurations, especially for large  $N$ . There exist parameterizations leading to good results (e.g., when  $N = 16$ ,  $r = 3$  for  $\kappa \leq 0.85$ , or  $r = 5$  for  $\kappa = 1$ ), but no single one performs systematically well across different



**Fig. 1.** Accuracy of different hierarchical structures ( $m = 2, 4, N$ , indicated by line type) for different values of  $\kappa \in \{0.75, 0.85, 1\}$  with varying number of options  $N$  and parameter  $r$  (line color). Other parameters:  $M = 100$ ,  $Q = 0.8$ ,  $\mathcal{T} = 1000$  time steps.






**Fig. 2.** Performance evaluation over 100 independent runs of different hierarchical structures ( $m = 2, 4, N$ ) with varying number of options  $N$  (point color) and parameter  $r$  (color intensity). For each value of  $N$ , the Pareto frontier is displayed connecting points that are non-dominated. Dominated points are smaller than non-dominated ones. Left:  $\kappa = 0.75$ . Right:  $\kappa = 0.85$ . Other parameters as in Fig. 1.

values of  $k$  and  $N$ . Similarly, the quad trees deployed for  $N = 16$  display a good overall performance for  $r = 3$ , but still not comparable with the binary trees.

Overall, we found that hierarchical decisions perform better as they allow to serialize the best-of- $N$  problem in a sequence of smaller problems, which can be better parameterized to deal with different complexity levels. However, high accuracy may come at the cost of slowing down the decision process, and this could be especially the case if a long sequence of decisions must be performed. In Fig. 2, we study the speed-accuracy trade-off for  $\kappa = 0.75$  (left) and  $\kappa = 0.85$  (right).<sup>1</sup> In both cases, the hierarchical approach produces solutions that lay on the Pareto frontier, often dominating non-hierarchical solutions. Hence, fast convergence is ensured also in case of a sequence of  $D$  decisions.

When the quality of the alternatives is not known a priori, a (slow) estimation is necessary from noisy observations. With hierarchical structures, estimation errors may lead to wrong decision in the early stages (e.g., nearer to the tree root) that cannot be easily recovered. Also, errors in quorum sensing can lead a whole group down the wrong path. In such conditions, the flat hierarchy could be advantaged. Our simulations demonstrate that estimation errors have an impact, but still hierarchical structures provide a sensible advantage (see Fig. 3 and 4). In this case,  $r = 0.5$  performs best also when  $\kappa = 1$ , because small differences between identical alternatives due to estimation errors get amplified, accelerating convergence towards any option. For  $r \geq 1$ , instead, estimation errors can lead to a wrong decision especially when alternatives are similar ( $\kappa = 0.85$ ). Flat trees and quad trees instead do not present solutions that are systematically good, similarly to the case in which agents could exploit perfect knowledge. The Pareto diagrams in Fig. 4 highlight that noisy estimation leads to a loss in accuracy and speed, especially with  $\kappa = 0.85$ . When  $N = 16$ , binary trees are not always Pareto optimal, but lay close to the frontier and can be accepted.

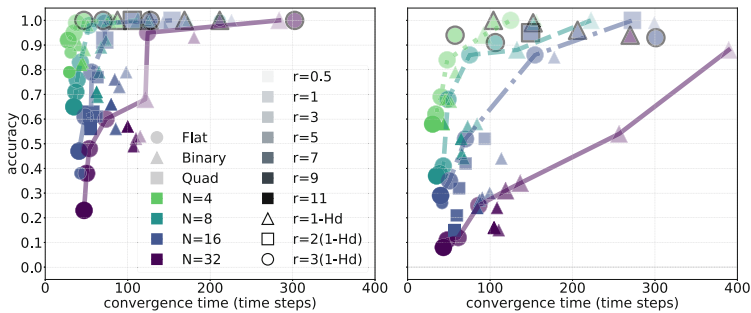
<sup>1</sup> For  $\kappa = 1$ , decision speed is very similar across different configurations, and the trade-off is dominated by solutions with high  $r$  that quickly converge to any option.



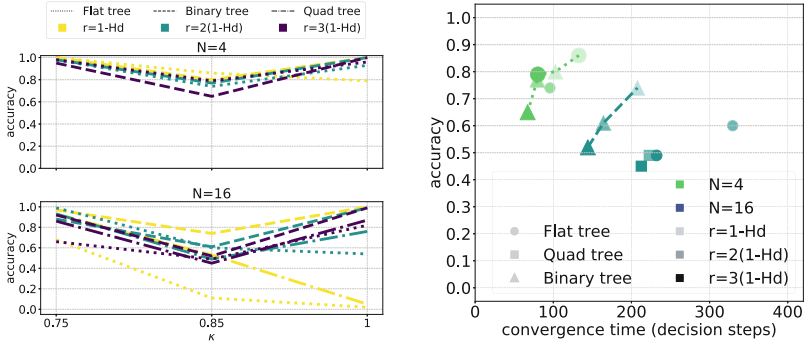
**Fig. 3.** Accuracy over 100 independent runs of different hierarchical structures with noisy observation and quorum sensing ( $M = 100$ ,  $Q = 0.8$ ,  $\lambda = 0.8$ ,  $\sigma_j^D = 1$ ,  $T_M = 12$ ,  $L_M = 10$ ,  $L_m = 2$ ,  $\mathcal{T} = 1000$  time steps). Black lines indicate the adaptive approach.

The slow evidence gathering process requires a small value of  $r$  to provide good accuracy, but this leads to a slow collective decision process. An adaptive approach can prove best if it gathers evidence to minimize uncertainties and increases speed when sufficient information is available. By linking the parameter  $r$  to the uncertainty, collective decisions can be both accurate and fast. Figure 3 and 4 show results of simulations with an adaptive approach, where the gain  $g$  has been tuned to maximize decision accuracy. An adaptive approach proves very advantageous especially with binary trees and complex decision problems, where both accuracy and speed are improved (see Fig. 4 right).

Finally, we analyse a proof of concept where mobile agents need to select the best area among  $N$  (see videos at [13]). Here, the adaptive approach is employed. Results shown in Fig. 5 demonstrate that binary trees provide the best solutions and also optimize the speed-accuracy trade-off. Hence, despite



**Fig. 4.** Pareto diagram corresponding to simulations with noisy observation and quorum sensing. Left:  $\kappa = 0.75$ . Right:  $\kappa = 0.85$ . Other parameters as in Fig. 3.



**Fig. 5.** Left: Accuracy of the spatial simulations for different configurations. Right: Pareto diagram for  $\kappa = 0.85$ . Other parameters:  $M = 100$ ,  $Q = 0.8$ ,  $\lambda = 0.8$ ,  $\sigma_j^D = 1$ ,  $T_M = 100$ ,  $L_M = 10$ ,  $L_m = 2$ ,  $T = 10000$  time steps.

the spatial correlations that may slow down decision-making, the hierarchical approach outperforms non-hierarchical configurations. Somewhat surprisingly, quad trees do not lead to good performance, despite they are the best choice for representing 2D spaces. Further studies should be performed to verify if different parameterizations lead to better results.

## 4 Conclusions

This study demonstrates that collective decision making can benefit from a hierarchical representation of the alternatives. A number of assumptions have been made, such as the a priori knowledge of  $N$  and of the tree structure. Such assumption can be relaxed, providing agents with the ability to build the hierarchy on the fly—possibly in a collective, decentralized way—or discovering it through observations, should this be related to the decision problem. Another assumption is related to the propagation of the maximum value from the leafs up the tree, which requires knowledge of what leaf is being observed from any non-leaf node. If such knowledge is not available, all observations must be aggregated at the non-leaf node, for instance via a (moving) average. This however reduces the ability to distinguish between different alternatives at the beginning of the process, as preliminary experiments have demonstrated (data not shown). We hypothesize that, by means of an adaptive sampling approach (e.g., Thompson sampling [20]), better alternatives could be sampled more frequently, leading to approximate the propagation of the maximum value. This will be studied in the future, along with implementation with Kilobots [19].

**Acknowledgments.** Vito Trianni acknowledges partial support from the project TAILOR (H2020-ICT-48 GA: 952215).

## References

1. Albani, D., Hönig, W., Nardi, D., Ayanian, N., Trianni, V.: Hierarchical task assignment and path finding with limited communication for robot swarms. *Appl. Sci.* **11**(7), 3115 (2021)
2. Albani, D., Manoni, T., Nardi, D., Trianni, V.: Dynamic UAV swarm deployment for non-uniform coverage. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018*, pp. 523–531 (2018)
3. Bareiss, D., van den Berg, J.: Generalized reciprocal collision avoidance. *Int. J. Robot. Res.* **34**(12), 1501–1514 (2015)
4. Bettstetter, C., Hartenstein, H., Pérez-Costa, X.: Stochastic properties of the random waypoint mobility model. *Wirel. Netw.* **10**(5), 555–567 (2004)
5. Bogacz, R.: Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* **11**(3), 118–125 (2007)
6. Brown, S., Steyvers, M., Wagenmakers, E.J.: Observing evidence accumulation during multi-alternative decisions. *J. Math. Psychol.* **53**(6), 453–462 (2009)
7. Busemeyer, J.R., Gluth, S., Rieskamp, J., Turner, B.M.: Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends Cogn. Sci.* **23**(3), 251–263 (2019)
8. Cassey, T.C., Evens, D.R., Bogacz, R., Marshall, J.A.R., Ludwig, C.J.H.: Adaptive sampling of information in perceptual decision-making. *PLoS ONE* **8**(11), e78993 (2013)
9. Gluth, S., Spektor, M.S., Rieskamp, J.: Value-based attentional capture affects multi-alternative decision making. *eLife* **7**, e39659 (2018)
10. Jin, L., Rückin, J., Kiss, S.H., Vidal-Calleja, T., Popović, M.: Adaptive-resolution field mapping using Gaussian process fusion with integral kernels. *IEEE Robot. Autom. Lett.* **7**(3), 7471–7478 (2022)
11. Kameda, T., Toyokawa, W., Tindale, R.S.: Information aggregation and collective intelligence beyond the wisdom of crowds. *Nat. Rev. Psychol.* **1**, 345–357 (2022)
12. Kaplan, E.L., Meier, P.: Nonparametric estimation from incomplete observations. *J. Am. Stat. Assoc.* **53**(282), 457–481 (1958)
13. Oddi, F., Cristofaro, A., Trianni, V.: Videos associated to “Best-of-N collective decisions on a hierarchy” (2022). <https://doi.org/10.5281/zenodo.6786939>
14. Pavlic, T.P., Hanson, J., Valentini, G., Walker, S.I., Pratt, S.C.: Quorum sensing without deliberation: biological inspiration for externalizing computation to physical spaces in multi-robot systems. *Swarm Intell.* **15**(1–2), 171–203 (2021)
15. Reina, A., Bose, T., Trianni, V., Marshall, J.A.R.: Psychophysical laws and the superorganism. *Sci. Rep.* **8**(1), 4387–8 (2018)
16. Reina, A., Marshall, J.A.R., Trianni, V., Bose, T.: Model of the best-of-N nest-site selection process in honeybees. *Phys. Rev. E* **95**(5), 052411–15 (2017)
17. Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., Trianni, V.: A design pattern for decentralised decision making. *PLoS ONE* **10**(10), e0140950–18 (2015)
18. Ross-Gillespie, A., Kümmerli, R.: Collective decision-making in microbes. *Front. Microbiol.* **5**(54) (2014)
19. Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., Nagpal, R.: Kilobot: a low cost robot with scalable operations designed for collective behaviors. *Robot. Auton. Syst.* **62**(7), 966–975 (2014)
20. Russo, D., Van Roy, B.: An information-theoretic analysis of Thompson Sampling. *J. Mach. Learn. Res.* **17**(1), 2442–2471 (2016)

21. Saha, A., Marshall, J.A.R., Reina, A.: Memory and communication efficient algorithm for decentralized counting of nodes in networks. *PLoS ONE* **16**(11), e0259736 (2021)
22. Sasaki, T., Pratt, S.C.: The psychology of superorganisms: collective decision making by insect societies. *Annu. Rev. Entomol.* **63**(1), 259–275 (2018)
23. Sridhar, V.H., et al.: The geometry of decision-making in individuals and collectives. *Proc. Natl. Acad. Sci. U.S.A.* **118**(50), e2102157118 (2021)
24. Valentini, G., Ferrante, E., Dorigo, M.: The best-of-N problem in robot swarms: formalization, state of the art, and novel perspectives. *Front. Robot. AI* **4**, 1–43 (2017)