Human-Collective Collaborative Target Selection

JASON R. CODY, United States Military Academy

KARINA A. ROUNDTREE and JULIE A. ADAMS, Oregon State University

Robotic collectives are composed of hundreds or thousands of distributed robots using local sensing and communication that encompass characteristics of biological spatial swarms, colonies, or a combination of both. Interactions between the individual entities can result in emergent collective behaviors. Human operators in future disaster response or military engagement scenarios are likely to deploy semi-autonomous collectives to gather information and execute tasks within a wide area, while reducing the exposure of personnel to danger. This article presents and evaluates two action selection models in an experiment consisting of a single human operator supervising four simulated collectives. The action selection models have two parts: (1) a best-of- n decision-making model that attempts to choose the highest-quality target from a set of n targets and (2) a quorum sensing task sequencing model that enables autonomous target site occupation. An original biologically inspired insect colony decision model is compared to a bias-reducing model that attempts to reduce environmental bias, which can negatively influence collective best-of- n decisions when poorer-quality targets are easier to evaluate than higher-quality targets. The collective decision-making models are compared in both supervised and unsupervised trials. The bias-reducing model without human supervision is slower than the original model but is 57% more accurate for decisions where evaluating the optimal target is more difficult. Human-collective teams using the bias-reducing model require less operator influence and achieve 25% higher accuracy with difficult decisions compared to the teams using the original model.

CCS Concepts: • **Computing methodologies** → **Multi-agent systems**; • **Human-centered computing** → *HCI theory, concepts and models*;

Additional Key Words and Phrases: Human-swarm interaction, collective decision making, swarm intelligence

ACM Reference format:

Jason R. Cody, Karina A. Roundtree, and Julie A. Adams. 2021. Human-Collective Collaborative Target Selection. *Trans. Hum.-Robot Interact.* 10, 2, Article 18 (March 2021), 29 pages.

<https://doi.org/10.1145/3442679>

This work was partially funded by the US Office of Naval Research Awards N000141210987, N00014161302, and N000141613025. The work of J. R. Cody was supported by the US Military Academy and the US Army Advanced Civil Schooling (ACS) program. The views and conclusions contained herein are those of the authors and are not to be interpreted as representing the official policies or endorsements, either expressed or implied, of the United States Military Academy, the US Army, the Department of Defense, or the US Government.

Authors' addresses: J. R. Cody, United States Military Academy, 606 Thayer Road, West Point, NY 10996; email: jason.cody@westpoint.edu; K. A. Roundtree and J. A. Adams, Oregon State University, 101 Covell Hall, Corvallis OR 97331; emails: roundtrk@oregonstate.edu, julie.a.adams@oregonstate.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

2573-9522/2021/03-ART18 \$15.00

<https://doi.org/10.1145/3442679>

1 INTRODUCTION

Collectives are large (i.e., hundreds or thousands of individuals) distributed groups of entities that use local sensing and communication to coordinate group activity [45, 54, 62]. The term *collective* refers to a robotic system that exhibits behaviors inspired by biological spatial swarms (e.g., fish schools [13]), colonies (e.g., ant colonies [29]), or a combination of both (e.g., honeybees [50]). Collective behaviors emerge from the decentralized interactions of their individual entities [45]. The collectives' behaviors are flexible in response to environmental changes, scale well with changes in population size, and are robust to individual entity failure [4]. Collectives are advantageous in large dangerous environments that lack communications infrastructure, which is common in military combat and disaster response scenarios [7, 21, 47]. A collective of inexpensive entities, supervised by an operator [5], can provide large area sensing, reduce personnel exposure to danger, and mitigate the effect of individual entity failure on the collective's performance [47, 62]. Collectives in military scenarios translate to hundreds, or thousands, of sensing platforms capable of synchronous task execution and surveillance across large areas that are difficult, dangerous, and expensive with human personnel [47]. Collectives in disaster response scenarios (e.g., in wildfires) provide a level of wide area persistent risk assessment, victim location, environmental sensing, and synchronous action (e.g., extinguishing fire) that is difficult to achieve without placing more personnel, or expensive uninhabited systems, in danger [7]. Semi-autonomous collectives that are capable of independent decisions (e.g., choose the best target) and simple task execution (e.g., occupy a location) are desirable in military engagements [21] and disaster response scenarios [7], because the operator and the collective are unlikely to have perfect information about the environment [3, 14]. Sharing decision-making responsibilities between the operator and the collective is advantageous but requires the operator to observe and influence the collective's decision-making process [14]. Further, single-operator supervision over one, or more, collectives requires that the operator's interaction tasks do not increase with collective population size, which can be accomplished with interaction strategies that enable the direct influence of collective behaviors [5, 62].

The collective behaviors in this article were inspired by the honeybee process of identifying, choosing, and emigrating to a new hive site. During honeybee emigration, a large group of honeybees depart the original hive exhibiting spatial swarm behaviors (similar to fish schooling [13]) and establish a temporary colony [50]. The honeybee scouts evaluate the quality of possible future hive sites, using the temporary location to pool information and facilitate a decision [54]. The honeybee scouts use a quorum detection strategy to determine when enough scouts support a single hive site, indicating a decision [50, 51]. Once the honeybee scouts choose a site, the honeybees transition back to a spatial swarm as they travel to the selected hive site [50]. Although certain ant colonies are also known for making similar best-of- n decisions [16], honeybees uniquely travel as a group between their temporary location and the chosen site.

Several decentralized robotic systems have attempted to adopt behaviors similar to honeybee behaviors [4, 42]. Individual entities in spatial swarms coordinate with neighboring individuals to cause the emergence of desirable group behavior (e.g., coordinated movement [18] or aggregation [2, 17]). Robotic colonies are inspired by insect colonies that make collective decisions including the allocation of individual entities to different tasks [12], or the selection of the best option from a discrete set of n available options, referred to as a collective best-of- n decision [42, 56]. A primary characteristic of colonies are decision-making hubs, where information is pooled to facilitate colony decisions [55]. Recent research has also identified the decision-making hub as a focal point for human interaction with collective decision-making processes [3, 14].

This article combines a collective best-of- n decision model with a quorum detection-based task sequencing model [37] to enable collectives to make decisions and execute actions associated with

those decisions. The resulting model, called an *explicit action selection* model, allowed simulated collectives to act as colonies while they evaluated targets from a stable location before transitioning to spatial swarm behaviors to occupy a target's site. Two existing challenges are addressed. First, features of the environment (e.g., target location) are known to affect the accuracy of the best-of- n decision-making model. The bias is positive when it improves the collective's ability to choose the highest-quality target and negative otherwise [55]. A bias-reducing best-of- n decision-making model introduced additional control mechanisms designed to reduce the influence of negative bias, which occurs when higher-quality targets are more difficult to discover and evaluate relative to nearby poorer-quality targets. The original model and the bias-reducing model were compared in an independent collective experiment to assess the collectives' decision making in a series of targeting decisions. The second challenge was creating an interface that enabled an operator to interact with the collectives as needed to confirm, influence, or cancel targeting decisions. The original and bias-reducing decision-making models were evaluated in an experiment where a single human operator supervised four collectives making a series of concurrent targeting decisions. The goal was to enable the operators to better inform collective decisions and provide the means to influence each collective's decision making.

2 RELATED WORK

Collective decision making and human interaction with decentralized robotic systems are active research areas (e.g., [27, 55]). Collective decision making includes two different types of decisions: task allocation and consensus achievement [4]. Allocation decisions divide a collective's individual entities between a finite set of tasks. Consensus achievement decisions require the collective to choose one option from a set of available options [54]. Continuous consensus decisions consider an infinite set of options (e.g., movement direction [18]), whereas discrete consensus decisions select one option from a finite set. Collective best-of- n decisions are discrete consensus decisions that select the highest-quality option of a finite set of n options [55].

Collective best-of- n decision strategies have enabled locally sensing groups of entities to determine the best resource (e.g., [6, 42, 53]), shortest path (e.g., [32, 40, 48, 49]), or future action (e.g., [37, 61]). A general opinion-based approach to collective best-of- n decision making requires the individual collective entities to maintain opinions favoring one of a finite set of n available options [54]. The collectives' individuals only disseminate their current favoring opinion within the collective decision-making hub, in a manner similar to emigration decisions made by *Apis mellifera* honeybees [34, 50] and *Temnothorax albipennis* ant [29] colonies. The individual collective entities periodically re-evaluate their currently supported option, which may be a specific location (e.g., [56, 57]) or a path (e.g., [32, 48]), by leaving the hub and performing an assessment. The dissemination of opinions within the hub is subject to *direct modulation*, as each individual collective entity shares its opinion with a frequency that is determined by the opinion's value. Features of the environment also cause the *indirect modulation* of individual opinion sharing. These environmental features cause an *environmental bias* that affects collective decision outcomes [54]. Environmental bias that increases opinion sharing frequency for optimal options is called *positive environmental bias*. *Negative environmental bias* occurs when the environment reduces the dissemination frequency for optimal sites and often causes poor decisions [55].

Individual collective entities alter their opinion based on an individual decision strategy [54]. Evaluated individual decision strategies include voter models (e.g., [40, 41, 56]), in which individuals randomly select a neighbor's opinion, and majority models (e.g., [32, 57, 61]), in which the individuals choose the majority opinion within a local neighborhood of individuals. One weakness of these strategies is an inability to choose the best option when that option changes during the decision-making process [39]. Recently, the introduction of a small number of stubborn agents,

or the introduction of a small probability of spontaneous opinion switching, have been shown to improve collective decisions in dynamic binary decision problems [39].

An opinion-based decision-making strategy, extended in this article, was inspired by *Apis mellifera* honeybee collectives when they choose and emigrate to one of an available set of future locations [42]. This strategy enables the collective to discover new options during deliberation between options and is value sensitive, which delays a decision when the choice is between poor-quality options [34, 42]. Reina et al. [42] demonstrated that by increasing the time between interactions, the collective was more resilient to negative environmental bias. Further efforts to improve accuracy despite negative environmental bias have included introducing additional individual entity control states based on distance from the decision-making hub [9, 11] and the modulation of agent interaction rates over time [53].

This work distinguishes between two types of collective action selection: implicit and explicit. Most collective best-of- n decision-making studies assume that a decision is complete when a consensus, or a majority opinion, has been reached [55]. The collectives make decisions *implicitly* by selecting an action, such as aggregating in a certain location, or foraging from a certain site, but do not require the collective's entities to respond to the collective's decision [55]. *Explicit* collective action selection requires the collective to make decisions and alter its behavior once a decision has been determined. Altering behavior is common when the collective must determine that a task is completed [37], or must move to a newly chosen location [10].

Explicit collective action selection is often achieved through the use of quorum detection mechanisms to halt deliberation during a best-of- n decision process, once a majority opinion is established [36]. When an individual collective entity detects that the number of its peers that shared a similar opinion exceeds a given threshold, it enters into a committed state and transitions the collective from option deliberation to a decision state [36]. A similar strategy enables the collective to execute a known sequence of tasks by making a series of decisions about whether the current task was completed [37]. A model combining quorum detection mechanisms of Parker and Zhang [37] with the best-of- n decision strategy of Reina et al. [42] enables a collective to stop deliberation between available options after a decision was reached, without reducing decision accuracy [10].

The best-of- n decision-making process has the potential to improve collective autonomy, but the process is imperfect. Human interaction with the collective decision-making process is necessary to correct for inaccuracies caused by environmental bias [55], influence collective decisions based on unique operator information [14], and speed up the decision process, which can be slow [10], particularly when deciding between options of similar quality [34]. Human influence is also needed to permit or cancel the execution of an action following a decision. Decision-making tasks are often the responsibility of the human operator in existing studies of human and collective interaction [27]. Studies of human control over spatial swarms that search (e.g., [18, 59, 60]), cover (e.g., [22, 23, 25]), or traverse (e.g., [33]) an area have identified that the performance of the human-robot team depends upon the human operator's observations, activities, and decisions [27]. Human interaction with colony behaviors, like task allocation and best-of- n decisions, offer an opportunity to share decisions between the operator and collective as a means of improving system fault tolerance, or its resilience to either human or colony error [14]. One collaborative task allocation study required a team of humans and simulated aerial vehicles to perform a set of tasks [12]. Human operators and robots updated a virtual pheromone map that informed individual robot behavior decisions from a set of security-related behaviors [12]. A honeybee-inspired shared decision-making strategy introduced fault tolerance in human interaction with a simulated colony's best-of- n site selection decisions [14]. A later user study with this strategy evaluated the use of different types of beacons placed in the environment to influence a colony's investigation and selection of sites [3]. The study examined the effect of site number, configuration, and the presence of hazards. The

Fig. 1. The original best-of- n process [42] (M_1 , left) and the bias-reducing extension of the model (M_2 , right). Individual collective entities in shaded states are latent and outside the hub. These entities explore (U^L) or evaluate (F_i^L) targets. Individual collective entities in white states are interactive inside the hub and may be recruited (U^I) or inhibited (F_i^I, F_i^D, F_i^R).

researchers argued that human influence over the decision process must be moderated to achieve desired team performance [3]. Future human-collective interactions will likely integrate the shared decision-making strategies applicable to colony behavior and the control strategies applicable to spatial swarms to develop semi-autonomous collectives [21]. Additional investigation is needed to determine how to improve collective best-of- n decision-making capabilities and how to effectively interact with one or more semi-autonomous collectives.

3 EXPLICIT COLLECTIVE ACTION SELECTION MODEL

The focal problem starts with the collective best-of- n decision problem, described in the second case study of Reina et al. [42]. A collective is required to explore an area and attempt to identify the highest-valued target within a search space. Each target, $i \in \{1, \dots, n\}$, had a value, $v_i \in (0, 1]$, and a location site, $loc_i \in \mathbb{R}^2$, that was a distance, d_i , from the collective's decision-making hub. Individual collective entities directly modulated their interactions based on the value, v_i , of the target they supported, which improved their ability to recruit other individuals to favor the same target. An accurate decision occurred when a majority of the individual collective entities favored the highest-valued target. Target distances from the collective decision-making hub caused a negative environmental bias when higher-quality targets were farther from the hub than lower-quality targets closer to the hub. The more distant higher-quality targets were less likely than the closer, poorer-quality targets to be selected, despite having higher relative values.

Expanding beyond the original problem, the presented model requires the collective to conduct explicit action selection to identify the highest-valued target, move the collective's decision-making hub from its original location to the chosen target's location, and initiate a subsequent decision. All evaluated models used a similar series of quorum detecting states to permit the collective to recognize when a decision was reached and move the collective. The two evaluated explicit collective action selection models differed according to each model's implementation of the collective best-of- n decision process, which are displayed in Figure 1. Individual collective entities interact within the hub, indicated by the white states, and do not interact outside the hub while assessing known targets or exploring, indicated by the shaded states.

The first model (M_1), as shown in Figure 1(a), implemented the best-of- n decision-making process of Reina et al. [42]. A population of N individual collective entities either favor a target i ($F_i^* \forall i \in \{1, \dots, n\}$) or are uncommitted to any target (U^*). Uncommitted latent individuals (U^L) explore outside the hub for targets and, with probability, P_{Y_i} , discover target i and return to the hub in an interactive favoring state (F_i^I). Uncommitted latent individuals eventually stop exploring, with probability P_U^I , and return to the hub in an uncommitted interactive state (U^I). Uncommitted interactive individuals (U^I) are recruited by interactive favoring individuals (F_i^I), with probability P_{ρ_i} , and enter the favoring latent state (F_i^L), or resume exploration (U^L), with probability P^L . Favoring latent individuals (F_i^L) assess target i before returning to the hub but may abandon their target with probability P_{α_i} and become uncommitted interactive (U^I). Favoring interactive individuals (F_i^I) are subject to being inhibited by individuals supporting targets other than i , with probability $P_{\sigma-i}$, which causes the inhibited individuals to enter the uncommitted latent state, U_L . Individual collective entities moved within the hub and while exploring outside the hub using correlated random walks [8]. Individuals moved directly between known targets and the hub. The random walks mixed the population within the hub, which frequently altered each individual's neighbors between that individual's opinion transmissions [42].

Environmental bias influences the original model (M_1) when target value is independent of target location [9, 11, 42]. Targets that are farther away from the decision-making hub are less likely to be chosen, because they have a lower probability of discovery (P_{Y_i}) and require more time for individual collective entities in state F_i^L to evaluate, which decreases P_i^I [11]. The second model, M_2 in Figure 1(b), extends model M_1 with two mechanisms that attempt to compensate for environmental bias: *interaction delay* and *interaction frequency modulation* [9, 11]. Interaction delay introduces two semi-interactive control states whose exit probabilities are determined by comparing the target's distance from the decision-making hub, d_i , to the minimum and maximum target distances (d_{min} and d_{max}) that define the collective's search space. After target discovery, individual collective entities enter state F_i^D that they leave with probability P_{D_i} , which is inversely proportional to the difference between the target's discovery time and the discovery time of a target located at d_{max} . Similarly, individuals that return from target reassessment enter state F_i^R , where they remain with probability P_{R_i} , which is inversely proportional to the difference between their current target's round trip time, τ_i , and the round trip time to a target at the maximum distance, τ_{max} . Interaction frequency modulation affects interaction rates between the individual collective entities.

The transition probabilities for both models depend on the target's value, v_i , and distance, d_i . The interactive transitions are related to the ratio of the individual entities supporting a target, N_i , to the total number of individual entities in the collective, N . Transition probabilities between latent and interactive states were determined using either the average round trip travel time, $\bar{\tau}$, or the actual round trip travel time to a target, τ_i . Model M_1 transition probabilities are described as follows:

$$P_{\alpha_i} = \alpha v_i^{-1}; \quad P_{Y_i} = v_i g(d_i); \quad g(d_i) = \frac{v_i \mu e^{\xi d_i}}{d_i}; \quad P_{\rho_i} = \frac{v_i N_i}{N}; \quad P_{\sigma-i} = \sum_{j \neq i} \frac{v_j N_j}{N}, \quad (1)$$

$$P_U^I = \bar{\tau}^{-1}; \quad P_i^I = \tau_i^{-1}; \quad P_L = 9P_U^I, \quad (2)$$

where α is the minimum abandonment rate, set to 5% per meter traveled [11, 42], and $g(d_i)$ is the estimated discovery probability, where d_i is the distance of target i from the collective hub and the values $\mu = 0.058$ and $\xi = -0.29$ were determined by fitting a curve to discovery data that was gathered from 50 simulations of the collective's exploration behavior [11, 42]. Equation (1) shows that abandonment probability (P_{α_i}) increases for lower-quality targets and that

higher-valued targets result in higher discovery (P_{Y_i}), recruitment (P_{ρ_i}), and inhibition ($P_{\sigma-i}$) rates. The transition probabilities for interaction delay and frequency modulation in model M_2 are described by

$$P_{D_i} = \frac{v_i}{g(d_{max})^{-1} - g(d_i)^{-1}}; \quad P_{R_i} = \frac{v_i}{\tau_{max} - \tau_i}, \quad (3)$$

$$M_i = \frac{d_i}{d_{min}}; \quad P_{\sigma-i}^M = \sum_{j \neq i} \frac{M_j v_j N_j}{N}; \quad P_{\rho_i}^M = \frac{M_i v_i N_i}{N}, \quad (4)$$

where Equation (3) defines the exit probabilities of states F_i^D and F_i^R as inversely proportional to the difference between a target at d_{max} and target i 's discovery and round trip times. The interaction frequency modulation equations (Equation (4)) are defined in terms of M_i , which is the ratio of the distance to target i from the decision-making hub and the minimum target distance. As d_i increases, the recruitment and inhibition attempt frequency for individual entities supporting target i also increases. Target distance is assumed to be less than d_{max} , which means $g(d_i) < g(d_{max})$ and $\tau_i < \tau_{max}$. The use of distance in these equations imposes a simplifying assumption that negative bias is consistent and known within the environment, which is only somewhat problematic in the context of aerial robots observing targets on the ground within the collective's range. The assumption is less realistic in ground or underwater robot scenarios, which may contain obstacles or unknown structured environments. Strategies for learning and responding to inconsistent and unknown negative bias is beyond the current work and left to future investigation.

Both models were extended with a series of quorum detecting states, similar to an existing task sequencing model [37], to permit the collectives to stop deliberation once a target had been chosen and move to the chosen target. Each individual collective entity maintained a queue of its last $k = 15$ interactions, similar to the original task sequencing model [37]. The queues were cleared each time an individual changed states [10]. Individuals detected a quorum after receiving k messages in a row from other individuals in the same state, supporting the same target. Queue messages were sent with probability $M_i P_U^L$. Upon detecting a quorum, favoring individuals (F_i^I) entered a committed state (C_i), remained in the hub, and caused individuals in the deliberation states to commit to target i upon interaction. Committed individuals transitioned to an action initiation state (I_i) after detecting a quorum, and interacted with individual collective entities in states F_i^I , U^I , and C_i to transition to the initiation state. Initiating individuals started execution (X_i) after detecting a quorum, or after waiting k message intervals without receiving a queue message. Executing individuals moved from the collective's current hub location to the location of the chosen target, i , and entered the completed, or done (D_i), state. Finally, individual collective entities in state D_i transitioned back to the uncommitted interactive state (U^I) after detecting a quorum and began the next decision from the new location.

4 COLLECTIVE INTERFACE

Three requirements guided the design of the human-collective interface for the action selection processes described in Section 3: collective state transparency [1], restricted communication with the decision-making hub [14], and flexible collective autonomy [12]. First, collective state transparency requires that the interface permits rapid observation of the collectives' decision-making processes [1]. Second, restricting communication to the collectives' decision-making hubs avoids assuming a global communication capability [14]. Finally, flexible collective autonomy implies that each collective is capable of making independent decisions [35], but that the operator adjusts the collective's autonomy using control mechanisms (e.g., the operator allows the collective to make an independent decision or takes over the decision process as needed). Behavior selection

Fig. 2. The human-collective interface mid-way through a scenario, showing the current locations of the four collectives and the locations of the discovered targets. Collective decision-making hubs are the boxes indicated by Roman numerals. Targets are the green and blue boxes indicated by integer identifiers. The top half of each target indicates the target's relative value (green), and the bottom half indicates the support of the highest-supporting collective (blue).

controls [12, 26], which permit the operator to transition individual collective entities between internal states, were used to simultaneously alter the collectives' level of autonomy and guide collective decision outcomes. The remainder of this section describes interface visualization and control mechanisms used in the *Collective Interface*, which is in the screenshot in Figure 2.

4.1 Interface Visualization

The Collective Interface enabled an operator to interact with multiple collectives simultaneously, and each collective was required to choose the highest-valued target within its search range. Collectives and discovered targets were geo-located on the map area in the center of the interface. A static central map provided ecological validity and emulated a task searching an urban environment for potential locations of interest. The map provided operator context, such as the environment's topography, which is necessary to accomplish real disaster response or military scenario missions. The Collective Requests area (lower left) enabled the operator to change individual collective entities' behavior states in a manner explained in Section 4.2. The Collective Assignments area (upper right) provided a summary of active and inactive requests sent to the respective collectives. Finally, the System Messages area (lower right) provided text alerts of collective activities (e.g., finding targets) and human input error feedback. Four collectives and 16 targets are shown on the simulator's central map in Figure 2.

Box symbols represented collectives and targets. Collectives were white boxes identified by a Roman numeral above four quadrants labeled for the individual collective entity states (as described in Section 3): *uncommitted* (U), *favoring* a target (F), *committed* to a target (C), or *executing* a move to a target (X). The executing state included collective individuals in the previously described initiating (I) and done (D) states. The opacity of each quadrant indicated the percentage

of the population in that state. Collective I, in Figure 2, is primarily uncommitted or favoring, as indicated by the largely opaque *U* and *F* quadrants. Targets were green and blue boxes, labeled in the top right with an integer value. The opacity of a target's green upper portion indicated the quality of the target, with brighter green indicating higher quality. The opacity of the blue area of each target indicated the highest percentage of favoring agents supporting that target. Target 0, above Collective I in Figure 2, is a high-valued target (bright green top) with some support (semi-transparent blue bottom). Newly discovered targets were initially transparent but displayed a green color as soon as two individual collective entities favored the target and returned to the hub. Early in the decision process, target values fluctuated as a result of the collective members' noisy estimates of target value. A target was highlighted with a blue outline when more than 30% of a collective's individuals supported it, indicating a collective was about to commit to a target (see Target 0 in Figure 2). The target was highlighted with a green outline as individual collective entities began to execute the movement to a new location (see Target 8 in Figure 2). The Collective II symbol indicates that the collective is executing a movement from its current location to Target 4. The movement of Collective II is represented by the green box, which moves from the collective's current position to the chosen target's location. At the end of the movement, the Collective II symbol replaces the Target 4 symbol.

Left clicking on a hub, or a target, selected it as the designated object of the operator's collective request (see Section 4.2). Left clicking on a hub also highlighted targets that were in range of the collective and were supported (white outline) or not supported (yellow outline). Right clicking on targets revealed an estimate of each collective's support for the target, as shown with Target 2. Right clicking on a hub revealed a detail flag, as shown for Collective III, with an estimate of the number of collective members in each of the four previously described states.

The abstract visualization enabled the operator to quickly identify the state of the collective's decision-making process. The operator supervised collective decisions using the Collective Interface. The operator influenced the decision-making process using available behavior selection controls.

4.2 Collective Controls

The operator adjusted the collective's autonomy by completing *investigate*, *abandon*, or *decide* requests using the Collective Request area located on the lower left-hand side of Figure 2. Operator activities that were not related to requests were considered observation actions. Recorded collective observation actions included determining which targets were in range of a particular collective and extra left clicks on targets. Completed operator requests were added to the Collective Assignments area shown at the top right of Figure 2. Erroneous requests were identified to prevent the operator from attempting to have collectives investigate targets outside their range, abandon targets that had not yet been evaluated, or decide for the collective prior to the collective exploring its search space. Erroneous requests were ignored and caused an error message to be displayed to the operator in the Systems Messages Area, which was located at the bottom right of the simulation.

The operator issued an *investigate* request to increase a collective's support for a specific target by transitioning 10 uncommitted entities (5% of the collective population) to the favoring state supporting the chosen target. Additional support for the same target was achieved by reissuing the investigate request. Investigate requests for targets outside the collective's range were unacceptable and reported to the operator as errors. The *abandon* request reduced a collective's support for a specific target by transitioning favoring individual entities in the decision-making hub to the uncommitted state. The abandon request only needed to be issued once for the collective to ignore a target. Abandon requests for targets with fewer than two favoring agents were considered erroneous and reported to the operator. The *decide* request committed two individual collective entities

(1% of the population) to the target designated by the operator. These individuals then committed other individuals through interactions that rapidly drove the collective to execute a movement to the operator's chosen target. Issuing a decide request before at least 30% of the collective's population favored that target resulted in an error message.

Issuing a request required the operator to select the desired request type from the drop down menu, select the desired collective and target, and select the commit button to execute the request (see lower right of Figure 2). The reset button cleared the request information without sending the request. Recently executed requests (e.g., Collective I: Abandon Target 3) were displayed in the upper right hand corner of the monitor area, in the Collective Assignments area. Green and red circles next to each request signified whether the request was in progress (green) or completed (red). Investigate request circles changed from green to red after 10 individual entities received and acknowledged the investigate request. Abandon requests remained active (green), once issued. Once a collective reached a decision, all prior requests associated with that particular collective were removed from the Collective Assignments area. Abandon requests were the only requests the operator was able to cancel, which required selecting the desired request in the Collective Assignments area and pushing the cancel assignment button.

5 EXPERIMENTAL DESIGN

Two experiments were conducted to test the following hypotheses. H_1 predicted that incorporating the bias-reducing mechanisms described in Section 3 increases the accuracy of collective best-of- n decisions. H_2 predicted that human-collective teams outperform the independent collectives in both action selection decision accuracy and decision time. Finally, H_3 predicted that the Collective Interface enables the human operator to understand and modify the collective action selection behavior. The first experiment was simulation based and compared the behavior of the original collective best-of- n decision process, model M_1 , to that of the bias-reducing decision process, model M_2 , without humans. The second experiment was human oriented and compared the performance of the two decision process models within human-collective teams. A third model was added to the second experiment, called M_3 , which served as a control and did not include an independent best-of- n decision process. Model M_3 is described in Section 5.2.

Both experiments involved four simulated robotic collectives that were required to make targeting decisions within an approximately 2 km² environment. Each collective had 200 simulated Unmanned Aerial Vehicles (UAVs) that were limited to a 500-m exploration radius from the collectives' central decision-making hubs. Simulated robots behaved according to the models described in Section 3, which enabled each collective to perform explicit action selection within its coverage area.

Similar trials were used in both experiments and were modeled after the problem described in Section 3. At the start of each trial, four collectives were positioned 600 m apart in a square formation, centered within the environment, which purposely overlapped the collectives' search spaces. Each collective's starting location had four targets that were randomly located at distances within the set {150, 250, 350} m. The values assigned to the targets were within the set {0.7, 0.8, 0.9, 1.0} and were assessed by individual collective entities with a 10% error. Target values were assigned based on initial decision difficulty. Easy initial decisions placed the optimal target value (1.0) closest to the collective decision-making hub (e.g., at distance of 150 m), whereas difficult initial decisions placed the optimal target value (1.0) farthest from the decision-making hub (e.g., at a distance of 350 m). Collectives were visible at the start of the trial and targets were revealed after at least one collective individual discovered it and returned to the collective's hub. Trials were randomly generated according to these specified target values and distances.

Each trial was divided into two sections: one with easy initial decisions and one with difficult initial decisions. Each section required at least six decisions (for a total of at least 12 decisions in the trial). Following the collective's first target selection and occupation of its chosen target site, the collective's next decision was either easy or difficult, depending upon the collective's location relative to the remaining targets. The collective hubs were reset to their initial locations before commencing the second trial section, and all targets from the prior trial section were removed.

A larger set of individual robot capabilities than those typically used in collective best-of- n decision evaluation were assumed (e.g., [36, 42, 56]). Simple ground platforms such as Kilobots [44] and e-Puck robots [31] have been optimized for laboratory research. Both types of robots operate for long periods of time between battery charges (e.g., 2 hours for the e-Puck and several hours for the Kilobot) and are capable of short-range (<25 cm) peer-to-peer communication. Each capability is central to many biologically inspired collective decision-making strategies (e.g., [19, 36, 42, 56]) but are difficult to recreate in UAV collectives. UAV peer-to-peer communication is complicated by the UAVs' high mobility [28, 46]. Many commercial rotor wing UAVs are capable of speeds of 60km/h but have limited flight times of approximately 25 minutes (e.g., [30]). Despite these challenges, the simulated agents were assumed to have emergent capabilities and moved at 60km/h for at least an hour of flight time and communicated via 30-m peer-to-peer transmissions (e.g., optical transmissions [24]). A simplifying assumption was that robots moved at constant speed, which is similar to prior research [10, 42], but notably deviates from practical UAV movement.

The remainder of this section describes the two experiments in detail. Section 5.1 describes the experiment and metrics used to compare the basic collective action selection model, model M_1 , to the bias-reducing model, model M_2 . Section 5.2 describes the methods and metrics used to compare models M_1 , M_2 , and M_3 as part of human-collective teams in human trials.

5.1 Experiment 1: Independent Collective Action Selection Models

The first experiment tested hypothesis H_1 by comparing the explicit action selection models described in Section 3 in a series of target selection decisions with no human influence. The primary independent variable was the original (M_1) or bias-reducing (M_2) collective action selection models. Each model was evaluated 10 times for 28 trials, as described at the beginning of this section. Secondary independent variables included the locations of the targets with respect to each collective's decision-making hub, the targets' values, and decision difficulty (e.g., easy or difficult). Each collective made 6 decisions in each trial section, for a total of 12 decisions per trial.

The dependent variables for comparing the models' performance were success rate and decision time. Success rate was the ratio of the correct number of decisions made by the collective to the total number of decisions. A collective made correct decisions by identifying and moving to the highest-valued target within its search area. Decision time was the time from the start of a collective decision to the completion of the collective's movement to its chosen target. The success rate and decision times were averaged over the 10 runs of each of the 28 trials to retain a similar sample size when comparing to the human performance in the second experiment.

5.2 Experiment 2: Collective Action Selection Models in Human-Collective Teams

The second experiment evaluated the two collective action selection models and a baseline model within human-collective teams using the Collective Interface Visualization. The primary and secondary independent variables were identical to the first experiment; however, a baseline model, called M_3 , was created. Unlike models M_1 (original) and M_2 (bias reduction), M_3 did not make independent decisions or deliberate between targets (e.g., recruitment and inhibition interactions were disabled). The human operator increased support for targets and made decisions for the M_3

collective using investigate and decide requests, respectively. The M_3 model established operators' baseline performance with a low autonomy collective decision-making algorithm.

Each operator completed three 20-minute trials, with each trial corresponding to one of the collective action selection models. The trials were assigned such that half the operators completed the evaluation with M_1 followed by M_2 , and the other half experienced the M_2 model first. Additionally, the trials were assigned such that half the operators for each model started the trial with difficult initial decisions and the other half began with easy initial decisions. All operators conducted the baseline M_3 trial last, with the easy initial decisions preceding the hard decisions. The baseline model was expected to benefit from the learning effects related to the problem, interface, and collective behavior when compared to the primary models. Each trial required 12 collective target selections to be made by the four collectives simultaneously, which ensured at least two collectives made two consecutive decisions in each trial section.

The dependent variables were success rate, decision time, request frequency, and intervention rate. Success rate and decision time were defined in Section 5.1. Request frequency was the number of requests, per decision, per minute issued by the human operators during the trials. The frequency of each request type (e.g., investigate, abandon, and decide) was also gathered. Intervention rate was the number of interventions for 12 decisions per operator. An intervention occurred when the collective had achieved at least 10% support for a target and the operator issued an abandon request. Interventions indicated that the operator overrode the collective's decision-making process for the collective to choose a different target.

Additional objective metrics included the operator actions associated with observing a collective's state (e.g., clicking on a collective or target to view state information). During the trials, operators were required to answer probe questions [20] used to determine their level of situational awareness [15]. Situational Awareness (SA) probe questions can determine a human operator's situational awareness during critical tasks [20] according to three levels of situational awareness of Endsley [15]: perception, comprehension, and projection. Perception (Level 1) questions determined the operators' ability to perceive the targets and collectives as well as attributes associated with each (e.g., "Which collective is investigating Target 1?"). Comprehension (Level 2) questions determined operators' understanding of perceived elements in relation to collective decisions (e.g., "Which target is the best choice for Collective III?"). Finally, projection (Level 3) questions determined operators' ability to estimate the collective's future state based on their perception and comprehension of the current state (e.g., "Which collective will make the next decision?"). Operators were asked four Level 1, five Level 2, and three Level 3 SA probes for each model. The SA probes were similar across operators but were written as templates and the specific collective or target was added during the experiment. The Level 1 question, "Which Collective is investigating Target _?", for example, was completed with an applicable target number during the trial, before asking the SA probe question.

Recorded subjective data included answers to a demographic questionnaire, performance on the Mental Rotations Test (MRT) of Vandenberg and Kuse [58], and a post-experiment questionnaire that rank ordered the models. Each trial ended with a post-trial questionnaire, the NASA Task Load Index (NASA-TLX) and a 3-D Situation Awareness Rating Technique (3-D SART) [52]. The NASA-TLX provided a workload estimate, which included the weighted summation of a variety of workload components including mental demand, effort, and frustration. The 3-D SART score for situational awareness was calculated using the perceived Situational Understanding (SU), Demands on Attentional Resources (DAR), and Supply of Attentional Resources (SAR), according to the following equation: $SART\ Score = SU - (DAR - SAR)$. The post-trial questionnaires focused on the perceived performance and the responsiveness of the collective during the trial. The

Table 1. Success Rate (%) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1 SIM	52.11 (39.41)	60 (0/100)	70.49 (35.29)	87.5 (0/100)	6.8 (17.85)	0 (0/100)
M_2 SIM	74.58 (18.39)	70 (20/100)	76.7 (26.87)	83.33 (0/100)	64.04 (26.38)	66.67 (0/100)

post-experiment questionnaire required the operators to rank order the different collective models according to responsiveness, performance, and ease of comprehension.

Twenty-eight participants from the Vanderbilt University campus and surrounding area completed the experiment. The 15 female and 13 male participants were predominantly in the 18- to 30-year age range, although 4 participants were between the ages of 31 and 50. The participants had at least completed high school. More than half the operators had completed (13) or were completing (11) an undergraduate degree. It is important to note that no gender effects were observed. Each operator began the experiment by completing the informed consent paperwork, the demographic questionnaire, and the MRT. Once these items were completed, the operators received a scripted introduction to the experiment and the simulator (5 minutes). Prior to each trial, the operators conducted 5-minute training sessions with the specific model that consisted of two collectives, with one collective required to make an easy decision and one required to make a difficult decision. Operators responded to the SA probes at increments of approximately one probe per minute during each trial. At the end of each trial, the operators completed the post-trial questionnaire, the NASA-TLX, and the 3-D SART. After all trials and post-trial data collection, the operators completed the post-experiment questionnaire.

6 RESULTS

The results of the independent model experiment are presented first and compared to the results of the human-collective team experiments to examine the operator's influence on each model's performance. Section 6.1 describes the performance of the collectives in the absence of human influence in terms of success rate and decision time. Section 6.2 presents the performance of the human-collective teams, compares the performance of the human-collective teams to the independent model collective experiment, analyzes the actions and awareness of the operators, and presents the subjective results from the post-trial and post-experimental questionnaires, the NASA-TLX and 3-D SART worksheets, and the MRT scores.

6.1 Experiment 1: Independent Collective Action Selection Models

The descriptive statistics for the independent collectives' performance according to success rate and decision time are provided in Tables 1 and 2, respectively. The *SIM* designation (e.g., M_1SIM and M_2SIM) is used to distinguish the independent model results, from the human-collective teaming results that use the same models. The bias-reducing model, M_2SIM , achieved higher accuracy with slower decision times when compared to the original model, M_1SIM , especially during difficult decisions. A Mann-Whitney-Wilcoxon test indicated significant effects between the models for overall success rate, $N = 672$, *degrees of freedom* (DOF) = 1, $U = 41,017$, $p < 0.001$, and success rate for difficult decisions, $N = 456$, $DOF = 1$, $U = 3,157.5$, $p < 0.001$. A Mann-Whitney-Wilcoxon test also indicated significant effects between the models for decision time: overall decisions,

Table 2. Decision Time (in Minutes) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1 SIM	2.46 (0.66)	2.44 (1.36/4.46)	2.18 (0.53)	2.14 (1.36/3.89)	3.2 (0.96)	3.05 (1.83/7.08)
M_2 SIM	4.79 (1.11)	4.79 (2.49/7.7)	4.17 (0.93)	4.1 (2.49/7.55)	5.77 (1.38)	5.62 (3.67/10.25)

$N = 672$, $DOF = 1$, $U = 4,320.5$, $\rho < 0.001$; easy decisions, $N = 578$, $DOF = 1$, $U = 1,851$, $\rho < 0.001$; and difficult decisions, $N = 456$, $DOF = 1$, $U = 2,617$, $\rho < 0.001$.

The results indicate that model M_1 SIM selected the best target consistently when that target was closer to the collective's decision-making hub than other poorer-quality targets. During difficult decisions, M_1 SIM selected a nearby poor-quality target rapidly, even after discovering the best target. The bias-reducing model, M_2 SIM, achieved a 57% higher accuracy than M_1 SIM for the difficult decisions, demonstrating significant resilience to negative environmental bias; however, the model was generally 2 minutes slower than M_1 SIM during easy decisions.

Further analysis revealed a strong negative correlation between decision time and success rate when using M_1 SIM for all decisions, $r = -0.54$, $\rho < 0.001$; a moderately negative correlation for easy decisions, $r = -0.37$, $\rho < 0.001$; and a moderately positive correlation for difficult decisions, $r = 0.44$, $\rho < 0.001$. A strong negative correlation between decision time and success rate was also observed for M_2 SIM for all decisions, $r = -0.57$, $\rho < 0.001$, and a moderately negative correlation for easy decisions, $r = -0.44$, $\rho < 0.001$. These correlations suggest that longer decision times tended to result in poor easy decisions by both models. During difficult decisions, the basic model, M_1 SIM, made just under 7% correct decisions but was more likely to be correct when the decision took longer.

6.2 Experiment 2: Collective Action Selection Models in Human-Collective Teams

The results of the human-collective team experiment are presented in four sections. Section 6.2.1 compares the success rates and decision times achieved between the performance of the independent collectives (M_1 SIM and M_2 SIM) and the human-collective teams (M_1 , M_2 , and M_3). Section 6.2.2 presents the observed human actions during the experiment with models M_1 and M_2 and the baseline model, M_3 . Section 6.2.3 presents the observation actions and responses to the SA probe questions. Finally, Section 6.2.4 presents the subjective results.

6.2.1 Comparison between the Independent Collectives and Human-Collective Team Experiments.

The success rate and decision time descriptive statistics for the human-collective teams are provided in Tables 3 and 4 [43]. A Kruskal-Wallis test across models M_1 SIM, M_2 SIM, M_1 , M_2 , and M_3 identified significant effects for success rate in all decisions, $\chi^2(4, N = 1,680) = 523.39$, $\rho < 0.001$; easy decisions, $\chi^2(4, N = 1,147) = 381.3$, $\rho < 0.001$; and difficult decisions, $\chi^2(4, N = 895) = 388.94$, $\rho < 0.001$. Significant effects for decision times were observed in all decisions, $\chi^2(4, N = 1,680) = 687.89$, $\rho < 0.001$; easy decisions, $\chi^2(4, N = 1,147) = 683.52$, $\rho < 0.001$; and difficult decisions $\chi^2(4, N = 895) = 376.4$, $\rho < 0.001$.

Pairwise comparisons between the independent collectives and human-collective teams indicated significant human influence. A Tukey and Kramer test revealed significant effects between M_1 SIM and M_1 for success rates in overall decisions ($\rho < 0.001$), easy decisions ($\rho < 0.001$), and difficult decisions ($\rho < 0.001$). A similar test also revealed significant effects between M_1 SIM and

Table 3. Success Rate (%) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	78.57 (41.09)	100 (0/100)	95.31 (21.19)	100 (0/100)	56.25 (49.78)	100 (0/100)
M_2	88.39 (32.08)	100 (0/100)	94.44 (22.97)	100 (0/100)	81.41 (39.03)	100 (0/100)
M_3	92.86 (25.79)	100 (0/100)	95.94 (19.79)	100 (0/100)	88.49 (32.03)	100 (0/100)

Table 4. Decision Time (in Minutes) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	3.01 (1.56)	2.48 (1.16/8.58)	2.1 (0.75)	1.88 (1.16/4.79)	4.22 (1.54)	4.03 (1.6/8.58)
M_2	3.97 (1.37)	3.64 (1.83/9.94)	3.37 (1.23)	3.09 (1.83/9.94)	4.67 (1.2)	4.57 (2.46/8.81)
M_3	5.32 (2.22)	4.78 (1.48/13.25)	4.67 (1.96)	4.2 (1.48/12.78)	6.24 (2.24)	6.03 (2.07/13.25)

M_1 for decision times in overall decisions ($\rho < 0.001$) and difficult decisions ($\rho < 0.001$), but not for easy decisions. The pairwise comparison between M_2SIM and M_2 identified significant effects for success rates in overall decisions ($\rho < 0.001$), easy decisions ($\rho < 0.001$), and difficult decisions ($\rho < 0.001$). Differences in the decision times between M_2SIM and M_2 were also significant in overall decisions ($\rho < 0.001$), easy decisions ($\rho < 0.001$), and difficult decisions ($\rho < 0.001$).

Pairwise comparisons also indicated that the collective action selection model significantly affected human-collective team performance. A Tukey and Kramer test revealed significant effects between M_1 and M_2 for success rates in overall decisions ($\rho = 0.04$) and difficult decisions ($\rho < 0.001$). Significant effects were also observed between M_1 and M_3 for success rates in overall decisions ($\rho < 0.001$) and difficult decisions ($\rho < 0.001$). No significant effects were observed for success rates between M_2 and M_3 for any decision category. A Tukey and Kramer test revealed significant effects for decision times between all human-collective results (M_1, M_2, M_3) in overall and easy decisions ($\rho < 0.001$ for each comparison). Significant effects were observed in pairwise comparisons between each of the collective action selection models (M_1 and M_2) and M_3 for decision times in difficult decisions ($\rho < 0.001$), but no significant effects were observed for the same metric between M_1 and M_2 in difficult decisions.

The operator's influence is evident in Figure 3, which compares the success rates and decision times of the independent models (M_1SIM and M_2SIM) to the operators teamed with the same models (M_1 and M_2) and the baseline, human-only, model (M_3). Operators improved the success rates of both collective models, although success rates were clearly higher for M_2 than M_1 when making difficult decisions, as shown on the right side of Figure 3(a). Human influence enabled a success rate more than 10 times higher than the independent collective, M_1SIM , for difficult decisions. Operators decreased decision times for both collective action selection models for easy

Fig. 3. A comparison of the success rate and decision times for the independent collectives, M_1SIM and M_2SIM , and the human-collective teams, M_1 (original), M_2 (bias reducing), and M_3 (baseline).

decisions, but the higher difficult decision accuracy achieved by M_1 also required the original model to be slowed down, as shown on the right-hand side of Figure 3(b). The operators using M_2 increased success rates while decreasing decision times in all decisions.

The teams with the bias-reducing model, M_2 , achieved approximately 9% higher accuracy overall and 25% higher accuracy during difficult decisions when compared to the M_1 teams. The M_2 teams required more than a minute longer for each easy decision than the M_1 teams, although the gap in decision times for difficult decisions was noticeably smaller (less than 30 seconds on average). A moderate positive correlation was found between the decision time and success rates for the M_1 teams when making difficult decisions, $r = 0.51$, $\rho < 0.001$. Weak correlations were observed for the M_1 teams in overall decisions, $r = -0.19$, $\rho < 0.001$; for the M_2 teams in overall, $r = -0.11$, $\rho = 0.05$, easy, $r = -0.18$, $\rho = 0.02$, and difficult decisions, $r = 0.18$, $\rho = 0.03$; and for the M_3 teams in difficult decisions, $r = 0.25$, $\rho < 0.01$. The analysis suggests that the human-collective teams improved difficult decision accuracy with longer decision times and that this relationship was most common in the M_1 teams during difficult decisions. The operators made the fastest decisions using M_1 but achieved the lowest accuracy in difficult decisions even after slowing the collective decision process down. The operators using the baseline model, M_3 , were the most accurate but also the slowest. Finally, the M_2 human-collective teams made more accurate decisions than the M_1 teams and made these decisions faster than operators with the M_3 teams.

6.2.2 Human Actions in the Human-Collective Team Experiment. The significant influence of operators, reported in Section 6.2.1, resulted from the actions and awareness of the operators themselves. This activity was critical to the success of the baseline model, M_3 , which required the human operators to control the collectives without the aid of a collective decision-making process. The results provided in Tables 3 and 4, as well as in Figure 3, show that the human operators achieved the highest success rates with the baseline model, although these rates were not found to be significantly greater than those of the M_2 human-collective team. The baseline decision times, however, were significantly longer than the other two models for all decisions.

The operator influenced the collectives' behavior directly by issuing the *investigate*, *abandon*, and *decide* requests described in Section 4.2. The descriptive statistics for the frequency of all

Table 5. Request Frequency (Number of Requests per Decision/Decision Time)

	Overall		Easy		Difficult	
	Mean	Median	Mean	Median	Mean	Median
	(SD)	(Min/Max)	(SD)	(Min/Max)	(SD)	(Min/Max)
M_1	0.57 (0.56)	0.48 (0/2.56)	0.61 (0.61)	0.54 (0/2.56)	0.5 (0.47)	0.4 (0/1.98)
M_2	0.66 (0.49)	0.55 (0/2.6)	0.7 (0.56)	0.6 (0/2.6)	0.6 (0.4)	0.51 (0/2.03)
M_3	1.33 (0.79)	1.18 (0.11/4.75)	1.4 (0.87)	1.21 (0.11/4.75)	1.22 (0.65)	1.08 (0.3/3.66)

Table 6. Number of Investigate Requests per Decision per Minute

	Overall		Easy		Difficult	
	Mean	Median	Mean	Median	Mean	Median
	(SD)	(Min/Max)	(SD)	(Min/Max)	(SD)	(Min/Max)
M_1	0.44 (0.49)	0.32 (0/2.56)	0.47 (0.54)	0.43 (0/2.56)	0.38 (0.42)	0.29 (0/1.65)
M_2	0.47 (0.44)	0.35 (0/2.27)	0.49 (0.5)	0.34 (0/2.27)	0.46 (0.37)	0.37 (0/1.63)
M_3	1.07 (0.71)	0.95 (0/4.3)	1.12 (0.79)	0.97 (0/4.3)	1 (0.58)	0.92 (0/2.81)

requests issued per decision per minute are summarized in Table 5. The operators issued requests least frequently with M_1 and slightly more often with the bias-reducing model, M_2 . Requests were most frequent with the baseline model, M_3 , which required consistent human control. A Kruskal-Wallis test revealed significant effects between the three models for request frequency for overall decisions, $\chi^2(2, N = 1,008) = 239.65, \rho < 0.001$; easy decisions, $\chi^2(2, N = 569) = 121.84, \rho < 0.001$; and difficult decisions, $\chi^2(2, N = 439) = 120.3, \rho < 0.001$. A positive correlation was found between request frequency and success rate with M_1 for all decisions, $r = 0.35, \rho < 0.001$, but the correlation was strongest for the difficult decisions, $r = 0.58, \rho < 0.001$. A weak negative correlation was found for M_3 during difficult decisions, $r = -0.19, \rho = 0.03$. These results suggest that M_1 was more likely to make accurate decisions when the operator issued additional requests to the collectives, especially during difficult decisions. Differences in operator request frequencies did not correspond to changes in success rates for the bias-reducing model, M_2 , or the baseline model, M_3 . The M_3 teams achieved the highest success rates but required twice as many requests and at least a minute longer per decision than either the basic or bias-reducing collective action selection models.

The type of request issued to the collectives characterizes the kind of influence the operators exerted. Tables 6, 7, and 8 provide the number of investigate, abandon, and decide requests issued per decision per minute. Investigate requests were the most common, given that the abandon and decide requests result in a persistent change to the collective's state and were required less frequently. A Kruskal-Wallis test identified significant effects for investigate requests in overall decisions, $\chi^2(2, N = 1,008) = 223.38, \rho < 0.001$; easy decisions, $\chi^2(2, N = 708) = 114.41, \rho < 0.001$; and difficult decisions, $\chi^2(2, N = 550) = 111.64, \rho < 0.001$. The success rates and investigate requests were weakly correlated for all decisions when using M_1 , $r = 0.36, \rho < 0.001$, and M_2 , $r = 0.11, \rho = 0.04$, as well as difficult decisions using M_3 , $r = -0.19, \rho = 0.03$. Investigate requests were strongly

Table 7. Number of Decide Requests per Decision per Minute

	Overall		Easy		Difficult	
	Mean	Median	Mean	Median	Mean	Median
	(SD)	(Min/Max)	(SD)	(Min/Max)	(SD)	(Min/Max)
M_1	0.1 (0.19)	0 (0/0.8)	0.12 (0.22)	0 (0/0.8)	0.08 (0.12)	0 (0/0.43)
M_2	0.15 (0.16)	0.14 (0/0.55)	0.19 (0.18)	0.22 (0/0.55)	0.1 (0.13)	0 (0/0.41)
M_3	0.23 (0.12)	0.21 (0/1.15)	0.26 (0.12)	0.24 (0.08/1.15)	0.2 (0.12)	0.17 (0/0.9)

Table 8. Number of Abandon Requests per Decision per Minute

	Overall		Easy		Difficult	
	Mean	Median	Mean	Median	Mean	Median
	(SD)	(Min/Max)	(SD)	(Min/Max)	(SD)	(Min/Max)
M_1	0.02 (0.09)	0 (0/0.6)	0.02 (0.08)	0 (0/0.6)	0.03 (0.1)	0 (0/0.5)
M_2	0.02 (0.07)	0 (0/0.38)	0.02 (0.07)	0 (0/0.38)	0.02 (0.07)	0 (0/0.3)
M_3	0.03 (0.07)	0 (0/0.36)	0.03 (0.07)	0 (0/0.36)	0.03 (0.07)	0 (0/0.35)

correlated with success rate for difficult decisions when using M_1 , $r = 0.62$, $p < 0.001$, indicating that greater human influence improved decision accuracy. The number of investigate requests was significantly higher for M_3 than for the other models, but changes in Investigate Request Frequencies for both M_2 and M_3 did not correspond with changes in success rates with either model for easy or difficult decisions.

Decide requests occurred most frequently with M_3 , as expected considering this model did not make independent decisions. A Kruskal-Wallis test identified significant effects between models for decide requests in overall decisions, $\chi^2(2, N = 1,008) = 164.34$, $p < 0.001$; easy decisions, $\chi^2(2, N = 569) = 80.24$, $p < 0.001$; and difficult decisions, $\chi^2(2, N = 439) = 66.07$, $p < 0.001$. A moderate positive correlation was observed for decide requests and success rate with M_1 in both overall decisions, $r = 0.2$, $p < 0.001$, and difficult decisions, $r = 0.36$, $p < 0.001$. A moderate, negative correlation was also observed for difficult decisions, $r = -0.21$, $p < 0.01$, with M_3 . The positive correlation for M_1 suggests that the operator often directed accurate decisions in difficult decisions.

Abandon requests were also less common than investigate requests, although a Kruskal-Wallis test revealed significant effects for abandon requests for overall, $\chi^2(2, N = 1,008) = 11.27$, $p < 0.01$, and easy decisions, $\chi^2(2, N = 569) = 14.10$, $p < 0.001$. A weak, negative correlation was observed between abandon request frequency and success rate for M_1 decisions overall, $r = -0.13$, $p = 0.02$, but no other correlations were identified.

The differences between the models with respect to abandon request frequency was less informative than the differences observed in the intervention rate, which captures the abandon requests that were used to force a collective to ignore a target that had already gained at least 10% of the population's support within the collective. The descriptive statistics for the number of interventions per operator are presented in Table 9. A Kruskal-Wallis test identified significant effects across the models for number of interventions in overall decisions, $\chi^2(2, N = 84) = 10.35$, $p < 0.01$.

Table 9. Number of Interventions (Abandoned Targets with 10% Support) per Participant

	Overall	
	Mean (SD)	Median (Min/Max)
M_1	4.79 (2.97)	4 (0/12)
M_2	2.21 (1.99)	1.5 (0/7)
M_3	5 (5.11)	3.5 (0/18)

Table 10. Collective Observations (%) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	77.98 (41.5)	100 (0/100)	65.62 (47.62)	100 (0/100)	94.44 (22.99)	100 (0/100)
M_2	86.01 (34.74)	100 (0/100)	80 (40.11)	100 (0/100)	92.95 (25.68)	100 (0/100)
M_3	90.18 (29.8)	100 (0/100)	88.32 (32.19)	100 (0/100)	92.81 (25.93)	100 (0/100)

These results indicate that although requests were less frequent with M_1 in general, the requests used were more likely to correspond to changes in the human-collective team's success rate than with the M_2 and M_3 models. The positive correlations between M_1 's success rate and the number of requests, number of investigation requests, number of interventions, and number of decide requests suggests that in difficult decisions, the operator pushed the collectives toward accurate decisions. Using available controls, the operators were able to take advantage of M_1 's speed during easy decisions, but were required to take control in difficult decisions to achieve a success rate similar to that achieved by M_2SIM , without human interaction, on difficult decisions. Further, the high success rate observed with the baseline model, M_3 , indicates that the operators were capable of driving the collectives to the correct decision, but these decisions took significantly longer to make and required many more operator actions to complete.

6.2.3 Situational Awareness. Operators using M_1 , M_2 , and M_3 maintained varying levels of situational awareness while interacting with the collectives. The descriptive statistics for operators' observation actions, including collective observations and target observations, are presented in Tables 10 and 11, respectively. The displayed percentages are the ratio of the number of collective and target observation actions over the total actions for each decision. The observation percentages differed significantly across the human-collective teams. A Kruskal-Wallis test identified significant effects between the models for collective observations in overall, $\chi^2(2, N = 1,008) = 19.95, p < 0.001$, and easy decisions, $\chi^2(2, N = 569) = 29.77, p < 0.001$. Weak negative correlations were found between success rate and collective observations in overall decisions for $M_1, r = -0.12, p = 0.03$, and $M_2, r = -0.12, p = 0.03$ (identical calculated values). A Kruskal-Wallis test identified significant effects between the models for target observations per decision by overall decisions, $\chi^2(2, N = 1,008)$

Table 11. Target Observations (%) per Decision

	Overall		Easy		Difficult	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	20.54 (40.46)	0 (0/100)	10.94 (31.29)	0 (0/100)	33.33 (47.3)	0 (0/100)
M_2	22.62 (41.9)	0 (0/100)	21.11 (40.92)	0 (0/100)	24.36 (43.06)	0 (0/100)
M_3	41.07 (49.27)	0 (0/100)	41.12 (49.33)	0 (0/100)	41.01 (49.36)	0 (0/100)

Table 12. Percentage of Correct SA Probe Questions by SA Level

	SA_0		SA_1		SA_2		SA_3	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	85.12 (17.18)	91.67 (8.33/100)	86.31 (21.06)	100 (0/100)	87.32 (18.23)	100 (25/100)	78.57 (31.38)	100 (0/100)
M_2	89.88 (10.96)	91.67 (58.33/100)	91.67 (11.11)	100 (66.67/100)	88.39 (14.6)	100 (60/100)	89.88 (20.46)	100 (33.33/100)
M_3	87.2 (10.75)	91.67 (58.33/100)	94.05 (13)	100 (66.67/100)	91.43 (12.68)	100 (60/100)	76.79 (16.57)	75 (50/100)

$= 42.47$, $\rho < 0.001$; easy decisions, $\chi^2(2, N = 569) = 49.38$, $\rho < 0.001$; and difficult decisions, $\chi^2(2, N = 439) = 9.31$, $\rho < 0.01$. Weak positive correlations were found between target observations and success rate with the M_3 teams for both overall decisions, $r = 0.14$, $\rho = 0.01$, and difficult decisions, $r = 0.16$, $\rho = 0.05$.

The operators made the most collective observation actions with the baseline model, M_3 , which was expected, although the number of collective observations during difficult decisions was similar for all of the models. Target observations were significantly more common with the baseline model, M_3 , across decision difficulty. Operators used fewer observation actions for easy decisions with M_1 when compared to M_2 , but this relationship was reversed for difficult decisions. The increased target observations for M_1 were not correlated with success rates for that model in difficult decisions, but they were more common, indicating that the operators were more likely to make several observations of M_1 in difficult decisions than they were for the same model in easier decisions. Target observations were notably more common with M_3 than the other models. As with the collective observations with M_1 , target observations were also significantly more common with M_1 for difficult decisions than the corresponding easy decisions.

The final objective metric observed were the operators' responses to the SA probe questions. The SA probe responses evaluated the operators' understanding of the scenario while interacting with the collectives. The descriptive statistics for the SA correct responses for each level of SA are provided in Table 12 [43]. No significant effects were observed between the models for overall SA probe response accuracy, which were above 85% overall. The response accuracy was lowest for SA_3 , which required the operators to forecast a model's future behavior (see Section 5.1). A Kruskal-Wallis test identified significant effects between the models for SA_3 probe questions, $\chi^2(2, N = 84) = 7.57$, $\rho = 0.02$. The operators achieved a significantly higher SA Level 3 correct

Table 13. SART Results for Overall Score, SU, DAR, and SAR with Overall = SU - (DAR - SAR)

	Overall Score		SU		DAR		SAR	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	6 (2.28)	6 (1/13)	5.75 (1)	6 (3/7)	5 (1.22)	5 (1/7)	5.25 (1.48)	5.5 (2/7)
M_2	6.68 (2.26)	6.5 (3/13)	6.07 (0.9)	6 (4/7)	5.07 (1.18)	5 (1/6)	5.68 (1.09)	6 (3/7)
M_3	6.39 (2.08)	6 (4/11)	6.07 (0.98)	6 (4/7)	5.04 (1.43)	5 (1/7)	5.36 (1.31)	5 (3/7)

response rate when using M_2 as compared to the other models. Across the SA levels within each model, the correct response rate differences were significant. Kruskal-Wallis tests also identified significant effects between levels for each model: M_1 , $\chi^2(2, N = 56) = 64.63$, $\rho < 0.001$; M_2 , $\chi^2(2, N = 56) = 73.68$, $\rho < 0.001$; and M_3 , $\chi^2(2, N = 56) = 52.337$, $\rho < 0.001$. The human-collective team correct response percentage dropped more than 8% between SA_1 and SA_3 when using M_1 and 18% for M_3 . The human-collective teams using M_2 , in contrast, experienced less than a 2% reduction between SA_1 and SA_3 . The SA probe correct response percentages demonstrate that the human operators maintained a consistent ability across all models to access information about the collectives (SA_1) and understand the collectives' processes (SA_2). The SA_3 correct response percentages suggest that the use of the M_2 model either improved the human operator's ability to forecast the collective state or afforded the human operator more opportunities to properly respond.

The situational awareness results demonstrate that the operators needed to take many more observation actions with M_3 than the collective action selection models. Across decision difficulty, the observation activity of the M_2 and M_3 teams were consistent, but the observation behavior for the M_1 team noticeably increased between easy and difficult decisions. The SA probe correct response rates were higher than anticipated for all three models, which indicates that the Collective Interface Visualization was sufficient to enable operator situational awareness. The consistently better SA_3 response percentages for human-collective teams using M_2 strongly suggests an advantage of this model over the other two for the human operator. Due to the simultaneous decision-making during the trials, associating individual SA probe questions to decision problem difficulty was not possible; however, these results suggest that future examination of SA probe responses under different decision difficulties is likely to further distinguish these models.

6.2.4 Subjective Results. This section presents the subjective data including the operator's reported situational awareness, workload, post-trial questionnaires, post-experiment questionnaires, and the MRTs. The descriptive statistics for the 3-D SART and NASA-TLX are presented in Tables 13 and 14 [43], respectively. A Kruskal-Wallis test did not reveal significant effects between the models for the overall 3-D SART score or any of the 3-D SART components. The scores were similar across the models but slightly higher for M_2 .

A Kruskal-Wallis test revealed no significant effects for the overall NASA-TLX score but identified significant effects between the models for mental demand, $\chi^2(2, N = 84) = 22.166$, $\rho < 0.001$, and temporal demand, $\chi^2(2, N = 84) = 8.8327$, $\rho = 0.012$. A pairwise comparison using a Tukey and Kramer test revealed significant effects between M_1 and M_3 ($\rho = 0.003$) and M_2 and M_3 ($\rho < 0.001$) for mental demand. A similar test revealed significant effects for temporal demand between M_1 and M_3 ($\rho = 0.023$) and between M_2 and M_3 ($\rho = 0.032$). The operators reported higher mental and temporal demand when using models M_1 and M_2 . Higher mental demand indicates higher

Table 14. Significant NASA-TLX Results for Overall Score, Mental Demand, and Temporal Demand

	Overall		Mental		Temporal	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	58.31 (17.63)	60.67 (9/89.33)	22.19 (6.38)	23.33 (1.00/31.67)	11.55 (8.41)	9.67 (0.00/28.33)
M_2	57.06 (16.47)	56.83 (5.67/83.33)	23.58 (6.28)	25.00 (3.00/31.67)	10.94 (7.60)	10.33 (0.00/24.00)
M_3	50.63 (17.56)	54.17 (9.33/80.33)	16.54 (9.10)	18.50 (0.00/33.33)	7.49 (6.57)	5.33 (0.00/22.67)

No significant effects observed for Physical, Performance, Effort, and Frustration NASA -TLX components (omitted).

Table 15. Post Trial Request Type Effectiveness Ranking (1 = Low, 7= High)

	Investigate		Abandon		Decide	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	5.07 (1.25)	5 (2/7)	6.04 (1.35)	6 (1/7)	5.71 (1.94)	7 (1/7)
M_2	4.75 (1.53)	5 (2/7)	6.18 (1.42)	7 (1/7)	5.57 (1.99)	6 (1/7)
M_3	5.18 (1.68)	5.5 (1/7)	5.29 (1.76)	5.5 (1/7)	6.54 (0.92)	7 (4/7)

required perceptual or decision-making activity. The higher mental demand reported for M_1 and M_2 suggests that operators experienced higher demand when sharing decision-making tasks with the other collectives. When making decisions, as in M_3 , the operators did not need to consider what the collective was doing. The higher temporal demand is likely due to the fact that both models, M_1 and M_2 , made independent decisions whether the human operator intervened or not. The use of these models likely introduced additional pressure on the human operator to move quickly to influence each collective's decisions before the collective made an independent decision. M_3 did not impose a similar pressure for the human operator to act.

The post-trial questionnaires required Likert scale responses, on a scale of 1 (worst) to 7 (best) regarding the effectiveness of each type of request, the collective model's responsiveness to requests, the collective model's independent target selection ability (performance), and ease of understanding (comprehension). The descriptive statistics are shown in Tables 15 and 16, respectively. A Kruskal-Wallis test identified significant effects between the models for the effectiveness of the abandon request, $\chi^2(2, N = 84) = 6.33, \rho = 0.04$, and the model's independent performance, $\chi^2(2, N = 84) = 6.8, \rho = 0.03$. The fact that the operators rated the effectiveness of the abandon request lowest for the baseline model is not surprising, since the model only persistently investigated targets dictated by the human. The low performance reported for M_1 is consistent with the lower performance of the independent model, M_1SIM , compared to M_2SIM and the lower performance of the M_1 human-collective teams compared to the M_2 human-collective teams.

The post-experiment questionnaire required the operators to rank order the different models according to overall responsiveness to requests, overall performance, and overall comprehension. The descriptive statistics for the post-experimental rankings are summarized in Table 17. A

Table 16. Post-Trial Performance and Understanding
Model Ranking (1 = Low, 7 = High)

	Performance		Understanding	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	5.11 (0.99)	5 (3/7)	5.39 (1.31)	5.5 (2/7)
M_2	5.54 (1.29)	6 (3/7)	5.82 (1.16)	6 (3/7)
M_3	5.75 (1.43)	6 (2/7)	5.93 (1.46)	7 (3/7)

Table 17. Post-Experiment Model Ranking (1 = Best, 3 = Worst)

	Responsiveness		Performance		Comprehension	
	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)	Mean (SD)	Median (Min/Max)
M_1	1.5 (0.51)	1.5 (1/2)	2 (1.02)	2 (1/3)	2.5 (0.51)	2.5 (2/3)
M_2	1.5 (0.51)	1.5 (1/2)	2 (1.02)	2 (1/3)	2.5 (0.51)	2.5 (2/3)
M_3	3 (0)	3 (3/3)	2 (0)	2 (2/2)	1 (0)	1 (1/1)

Kruskal-Wallis test revealed significant effects for responsiveness, $\chi^2(2, N = 28) = 62.25, \rho < 0.001$, and comprehension, $\chi^2(2, N = 28) = 62.25, \rho < 0.001$ (identical values). The ranking of the baseline model, M_3 overall is interesting, as it was ranked consistently the lowest for responsiveness, the highest for comprehension, and between the remaining models for performance.

The final subjective metric was the MRT scores. The MRT has a minimum score of 0 and a maximum possible score of 24. The operators' MRT scores resulted in a 10.9 mean score (standard deviation = ± 5.5 , median = 10, minimum = 1, maximum = 24) [43]. These results were virtually identical to the results of a large study comprised of 636 operators with a 10.8 mean score and a standard deviation of ± 5 [38]. The results of the MRT were compared to the human-collective team success rates for each problem type. A Spearman correlations test only identified weak to moderate positive correlations between the MRT results and success rates with model M_1 for overall decisions, $r = 0.2, \rho < 0.001$; easy decisions, $r = 0.15, \rho = 0.03$; and difficult decisions, $r = 0.36, \rho < 0.001$. These findings suggest that operators with higher spatial awareness were slightly more likely to make accurate decisions with the M_1 collective action selection model.

7 DISCUSSION

The results of the experiments supported the overall research goal. The bias-reducing mechanisms enabled the collective to make better decisions with, or without, a human operator when compared to the model without these mechanisms. The Collective Interface Visualization and interaction strategy were sufficient for the human operator to maintain awareness of the collectives' states and influence the best-of- n decision process to improve the human-collective team's decision accuracy and time. When using the bias-reducing mechanisms, the interaction strategy demonstrated

significant fault tolerance [14], which means the team's success was resilient to changes in human operator behavior. The reliance of the human-collective team on the actions of the human operator differed depending upon whether or not the collective action selection model used bias-reducing mechanisms. When using model M_1 , the success of the human-collective team required the human operator to take control during difficult decisions to overcome the model's susceptibility to negative environmental bias. When using M_2 , the human-collective team made accurate decisions with relative indifference to the frequency of human operator actions.

The independent collective experiment supported hypothesis H_1 , which predicted that the bias-reducing mechanisms of M_2SIM improve decision accuracy when compared to the original model, M_1SIM . The results provide support of the well-known trade-off between speed and accuracy in these types of decisions [57]. The original model is fast when making decisions in which the environmental bias favors the best target (e.g., the best target is closest or equidistant to other targets). When required to select the best target, regardless of the target's distance from a collective's decision-making hub, M_1SIM is unreliable and makes poor decisions in difficult decision problems. M_2SIM 's resistance to environmental bias is a significant advantage over the original model, especially when accuracy in the best-of- n decision is more important than the time required to make a decision. Future investigations will determine if other bias-reducing mechanisms reduce the decision times achieved by M_2SIM (e.g., modulating interaction rates over time rather than distance [53]), or if collectives are capable of adjusting to unknown dynamic environmental bias (e.g., when unknown obstacles, or structures impede the movement of individual collective entities).

Hypothesis H_2 , which predicted that human influence improves the decision processes of the collective action selection models, was supported. The human-collective teams (M_1 and M_2) made faster, more accurate decisions than the independent collectives (M_1SIM and M_2SIM). The human-collective teams using the baseline model, M_3 , were very accurate but required the most human operator actions and made the slowest decisions of all human-collective teams. Comparing the behavior of the M_1 and M_2 human-collective teams reveals an important distinction between the two collective action selection models. Unlike the human-collective teams using M_2 , the M_1 teams' success rates were more strongly correlated with the human operators' request frequencies, use of investigate requests, use of decide requests, and spatial awareness (determined using MRT results). Further, the human operators were less than half as likely to intervene in the collective's decisions when using M_2 than when they were using M_1 , which may have occurred to correct M_1 's faster but less accurate decision making. The lack of similar observations with the M_2 human-collective teams supports that the M_2 model is significantly more tolerant to changes in human operator behavior than the M_1 model. Despite the objective results, the operators struggled to distinguish the models subjectively, even though they corrected the collectives' decision-making process more frequently with M_1 and achieved higher decision accuracy with M_2 .

The final hypothesis, H_3 , which stated that the Collective Interface Visualization and interaction strategy were sufficient to enable human interaction with the collective action selection models, was also supported. The operators' influence increased the selection of the highest-valued targets when compared to the collective decision processes alone (M_1SIM and M_2SIM). The operators' use of investigate and abandon requests to slow down the M_1 decisions and speed up the M_2 decisions demonstrates operator control over both collective action selection models to improve the human-collective team's decision accuracy and decision time. The operators responded consistently well to the SA_1 (perception) and SA_2 (comprehension) probe questions, suggesting that the operators perceived and understood the collectives' states. The ability to forecast a collective's future state (SA_3) was significantly lower for the human-collective teams when using models M_1 and M_3 . Operators may have been overloaded with the rapid change of collective behaviors and necessary adjustments to improve decision accuracy when using M_1 , whereas M_3 may have preoccupied

operators with the amount of influence required in the decision-making process. When using M_2 , the human operators were either afforded greater opportunity or capability to answer the SA_3 questions correctly. The objective observations contrast with the reported NASA-TLX workload, in which operators reported relatively equal mental and temporal demand between M_1 and M_2 , but the lowest demands with M_3 . This contrast between the objective and subjective measures is partially explained for the baseline model by the experimental design. Operator experience during the two trials preceding the M_3 trial likely positively influenced some of the subjective measures, including workload and model comparisons, due to learning affects. Additionally, since human operators controlled collective decision-making states when using model M_3 , they were expected to better understand the model's state. A challenge imposed by the experimental design was the inability to categorize subjective measures between difficult and easy decisions. Most of the objective measures differed between the decision difficulty, but the number of interventions, as well as the NASA-TLX, 3-D SART, and operator evaluations, were not assessed by decision difficulty.

Quantifying the Collective Interface Visualization's effectiveness requires its comparison with other visualization strategies. The Collective Interface Visualization enabled the operators to interact with each collective as a single entity rather than as organizations of 200 independent entities. Operators often observed a collective's decision-making state by reviewing the percentage of collective individuals performing each available behavior and the support visible in each collective's considered targets. An operator's requests to a collective's individuals were limited to the collective hub, with a report confirming to the operator that a request had been issued properly. The results demonstrate that operators consistently, and effectively, observed the collective state and influenced the collectives' decision-making processes.

Some design limitations of the Collective Interface Visualization can be improved to promote better human-collective behaviors. The inability to observe individual collective entity activities and the immediate changes in collective behavior after executing particular requests may have caused operators to reissue requests. The messages and status indicators (i.e., circular colored icons) provided in the Collective Assignments area were intended to communicate information about operator requests. The information may have not been salient enough to attract the operator's attention to the Collective Assignments area to confirm that their requests were executed or still in progress. Dividing the Collective Assignments area into two sections, one explicitly for executed requests and another for ongoing requests, instead of interpreting the colored icons, may help improve operator comprehension of a request's status. Alternatively, the words "executed" highlighted in red, or "in progress" highlighted in green, can be added to each request line drawing attention through the use of color and providing redundancy through text. The effectiveness of these design changes are only useful assuming the operators are not colorblind, as was done in this evaluation. The difference between the collective action selection models may be more perceivable if directional information, in the form of arrows or streamlines emitting from the collective's hub, indicated which targets were being favored. Operators may be able to see the bias toward nearby targets visually using the M_1 model versus the consistent favoring of high-valued targets, irrespective of the distance to the hub, using the M_2 model. Future evaluations are needed to affirm whether design changes to the Collective Interface Visualization improve human-collective behaviors, or if alternative visualization strategies are more effective.

Despite these limitations and future recommendations, the results are promising. The presented bias-reducing mechanisms embedded into the M_2 model enabled more accurate collective decisions. Human-collective teams using the bias-reducing model were more accurate and relied less on human influence than the teams that did not use these models. Finally, an abstract visualization that aggregates collective state information enabled the human operators to successfully supervise and influence the decision making of multiple collectives simultaneously.

8 CONCLUSION

This work evaluated a human-collective strategy intended to improve the collective's ability to inform human operator decisions and the human operator's ability to supervise and influence the collective's behavior. Two novel collective action selection models were evaluated to determine the effectiveness of new bias-reducing mechanisms on the collectives' decisions, the ability of human operators to improve the collectives' decision making, and the feasibility of a novel human-collective interface. The bias-reducing model made more accurate, but slower, collective best-of- n decisions than the model without these mechanisms in the absence of human operator influence. Human influence improved the collective decision accuracy and reduced decision times over the collectives alone for both action selection models. The collective without the bias reduction required human operator influence to drive the human-collective team to make accurate decisions. Further, the evaluated Collective Interface was sufficient to enable a human operator to influence four collectives, comprised of 200 individuals each, simultaneously solving the sequential best-of- n problem in static targeting scenarios. Although promising, this article highlights the need for future work to examine different visualization strategies, alternative bias-reducing mechanisms, or other human-collective control mechanisms to realize the benefits of human-collective shared decision making in real-world environments.

REFERENCES

- [1] Julie A. Adams, Jessie Y. C. Chen, and Michael A. Goodrich. 2018. Swarm transparency. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI'18)*. ACM, New York, NY, 45–46. DOI: <https://doi.org/10.1145/3173386.3177008>
- [2] Farshad Arvin, Ali Emre Turgut, Farhad Bazyari, Kutluk Bilge Arikan, Nicola Bellotto, and Shigang Yue. 2014. Cue-based aggregation with a mobile robot swarm: A novel fuzzy-based method. *Adaptive Behavior* 22, 3 (2014), 189–206. DOI: <https://doi.org/10.1177/1059712314528009>
- [3] C. C. Ashcraft, M. A. Goodrich, and J. W. Crandall. 2019. Moderating operator influence in human-swarm systems. In *Proceedings of the 2019 IEEE International Conference on Systems, Man, and Cybernetics (SMC'19)*. 4275–4282.
- [4] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. 2013. Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence* 7, 1 (2013), 1–41.
- [5] Daniel S. Brown, Michael A. Goodrich, Shin-Young Jung, and Sean Kerman. 2016. Two invariants of human-swarm interaction. *Journal of Human Robot Interaction* 5, 1 (2016), 1–31. DOI: <https://doi.org/10.5898/JHRI.5.1.Brown>
- [6] Alexandre Campo, Simon Garnier, Olivier Dédriche, Mouhcine Zekkri, and Marco Dorigo. 2011. Self-organized discrimination of resources. *PLoS One* 6, 5 (May 2011), 1–7. DOI: <https://doi.org/10.1371/journal.pone.0019888>
- [7] Daniel Carrillo-Zapata, Emma Milner, Julian Hird, Georgios Tzoumas, Paul J. Vardanega, Mahesh Sooriyabandara, Manuel Giuliani, Alan F. T. Winfield, and Sabine Hauert. 2020. Mutual shaping in swarm robotics: User studies in fire and rescue, storage organization, and bridge inspection. *Frontiers in Robotics and AI* 7 (2020), 53. DOI: <https://doi.org/10.3389/frobt.2020.00053>
- [8] Edward A. Codling, Michael J. Plank, and Simon Benhamou. 2008. Random walk models in biology. *Journal of the Royal Society Interface* 5, 25 (2008), 813–834. DOI: <https://doi.org/10.1098/rsif.2008.0014>
- [9] Jason R. Cody. 2018. *Discrete Consensus Decisions in Human-Collective Teams*. Ph.D. Dissertation. Vanderbilt University, Nashville, TN.
- [10] Jason R. Cody and Julie A. Adams. 2017. An evaluation of quorum sensing mechanisms in collective value-sensitive site selection. In *Proceedings of the 2017 International Symposium on Multi-Robot and Multi-Agent Systems (MRS'17)*. 40–47. DOI: <https://doi.org/10.1109/MRS.2017.8250929>
- [11] Jason R. Cody and Julie A. Adams. 2017. Reducing distance bias in value-sensitive swarm site selection. In *Proceedings of the International Symposium on Swarm Behavior and Bio-Inspired Robots*. 223–230.
- [12] Giles Coppin and François Legras. 2012. Autonomy spectrum and performance perception issues in swarm supervisory control. *Proceedings of the IEEE* 100, 3 (March 2012), 590–603.
- [13] Iain D. Couzin, Jens Krause, Nigel R. Franks, and Simon A. Levin. 2005. Effective leadership and decision-making in animal groups on the move. *Nature* 433, 1 (2005), 513–516.
- [14] Jacob W. Crandall, Nathan Anderson, Chace Ashcraft, John Grosh, Jonah Henderson, Joshua McClellan, Aadesh Neupane, and Michael A. Goodrich. 2017. Human-swarm interaction as shared control: Achieving flexible

- fault-tolerant systems. In *Engineering Psychology and Cognitive Ergonomics: Performance, Emotion and Situation Awareness*, Don Harris (Ed.). Springer International, 266–284. DOI : https://doi.org/10.1007/978-3-319-58472-0_21
- [15] Mica R. Endsley. 1995. Toward a theory of situation awareness in dynamic systems. *Human Factors* 37, 1 (1995), 32–64. DOI : <https://doi.org/10.1518/001872095779049543>
- [16] Nigel R. Franks, Stephen C. Pratt, Eamon B. Mallon, Nicholas F. Britton, and David J. T. Sumpter. 2002. Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357, 1427 (2002), 1567–1583.
- [17] Simon Garnier, Jacques Gautrais, Masoud Asadpour, Christian Jost, and Guy Theraulaz. 2009. Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior* 17, 2 (2009), 109–133. DOI : <https://doi.org/10.1177/1059712309103430>
- [18] Michael A. Goodrich, P. B. Sujit, Sean Kerman, Brian Pendleton, and José Pinto. 2011. *Enabling Human Interaction with Bio-Inspired Robot Teams: Topologies, Leaders, Predators, and Stakeholders*. Technical Report 2011-1. Human Centered Machine Intelligence Laboratory, Brigham Young University, Provo, UT.
- [19] Musad Haque, Christopher Ren, Electa Baker, Douglas Kirkpatrick, and Julie A. Adams. 2016. Analysis of swarm communication models. In *Proceedings of the International Workshop on Combinations of Intelligent Methods and Applications*. 29–37.
- [20] Curtis M. Humphrey. 2009. *Information Abstraction Visualization for Human-Robot Interaction*. Ph.D. Dissertation. Vanderbilt University, Nashville, TN.
- [21] Jules Hurst. 2017. Robotic swarms in offensive maneuver. *Joint Force Quarterly* 87 (Oct. 2017), Article 1326017.
- [22] Shin-Young Jung, Daniel S. Brown, and Michael A. Goodrich. 2013. Shaping Cousin-like torus swarms through coordinated mediation. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*. 1834–1839.
- [23] Shin-Young Jung and Michael A. Goodrich. 2013. Multi-robot perimeter-shaping through mediator-based swarm control. In *Proceedings of the International Conference on Advanced Robotics*. 1–6.
- [24] Asaad Kaadan, Hazem H. Refai, and Peter G. LoPresti. 2014. Multielement FSO transceivers alignment for inter-UAV communications. *Journal of Lightwave Technology* 32, 24 (Dec. 2014), 4785–4795. DOI : <https://doi.org/10.1109/JLT.2014.2364795>
- [25] Zsolt Kira and Mitchell A. Potter. 2009. Exerting human control over decentralized robot swarms. In *Proceedings of the International Conference on Autonomous Robots*. 566–571.
- [26] Andreas Kolling, Katia Sycara, Steven Nunnally, and Michael Lewis. 2013. Human-swarm interaction: An experimental study of two types of interaction with foraging swarms. *Journal of Human Robot Interaction* 2, 2 (2013), 103–129. DOI : <https://doi.org/10.5898/JHRI.2.2.Kolling>
- [27] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. 2016. Human interaction with robot swarms: A survey. *IEEE Transactions on Human-Machine Systems* 46, 1 (Feb. 2016), 9–26. DOI : <https://doi.org/10.1109/THMS.2015.2480801>
- [28] Arun K. Majumdar. 2015. *Free-Space Optical (FSO) Platforms: Unmanned Aerial Vehicle (UAV) and Mobile*. Springer, New York, NY, 203–225. DOI : https://doi.org/10.1007/978-1-4939-0918-6_6
- [29] Eamonn B. Mallon, Stephen C. Pratt, and Nigel R. Franks. 2001. Individual and collective decision-making during nest site selection by the ant *Leptothorax albigipennis*. *Behavioral Ecology and Sociobiology* 50, 4 (2001), 352–359.
- [30] Drone Business Marketer. 2017. 13 Best Drones with the Longest Flight Time—2018 Buyer’s Guide. Retrieved January 28, 2021 from <https://dronebusinessmarketer.com/best-drones-with-longest-flight-time/>.
- [31] Francesco Mondada, Michael Bonani, Xavier Raemy, James Pugh, Christopher Cianci, Adam Klapotcz, Stephane Magnenat, Jean-Christophe Zufferey, Dario Floreano, and Alcherio Martinoli. 2009. The e-puck, a robot designed for education in engineering. In *Proceedings of the Conference on Autonomous Robot Systems and Competitions*, Paulo J. S. Gonçalves, Paulo J. D. Torres, and Carlos M. O. Alves (Eds.), Vol. 1. IPCB: Instituto Politécnico de Castelo Branco, Castelo Branco, Portugal, 59–65. <http://www.est.ipcb.pt/robotica2009/>.
- [32] Marco A. Montes de Oca, Eliseo Ferrante, Alexander Scheidler, Carlo Pinciroli, Mauro Birattari, and Marco Dorigo. 2011. Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence* 5, 3–4 (2011), 305–327.
- [33] Sasanka Nagavalli, Nilanjan Chakraborty, and Katia Sycara. 2017. Automated sequencing of swarm behaviors for supervisory control of robotic swarms. In *Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA’17)*.
- [34] Darren Pais, Patrick M. Hogan, Thomas Schlegel, Nigel R. Franks, Naomi E. Leonard, and James A. R. Marshall. 2013. A mechanism for value-sensitive decision-making. *PLoS One* 8, 9 (2013), 1–9. DOI : <https://doi.org/10.1371/journal.pone.0073216>
- [35] Raja Parasuraman, Thomas B. Sheridan, and Christopher D. Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 30, 3 (May 2000), 286–297. DOI : <https://doi.org/10.1109/3468.844354>

- [36] Chris A. C. Parker and Hong Zhang. 2009. Cooperative decision-making in decentralized multiple-robot systems: The best-of-N problem. *IEEE/ASME Transactions on Mechatronics* 14, 2 (2009), 240–251.
- [37] Chris A. C. Parker and Hong Zhang. 2010. Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence* 4, 3 (2010), 199–220.
- [38] Michael Peters, Bruno Laeng, Kerry Latham, Marla Jackson, Raghad Zaiyouna, and Chris Richardson. 1995. A redrawn Vandenberg and Kuse mental rotations test—different versions and factors that affect performance. *Brain and Cognition* 28, 1 (1995), 39–58.
- [39] Judhi Prasetyo, Giulia De Masi, and Eliseo Ferrante. 2019. Collective decision making in dynamic environments. *Swarm Intelligence* 13, 3–4 (2019), 217–243. DOI: <https://doi.org/10.1007/s11721-019-00169-8>
- [40] Andreagiovanni Reina, Marco Dorigo, and Vito Trianni. 2014. Towards a cognitive design pattern for collective decision-making. In *Swarm Intelligence. Lecture Notes in Computer Science*, Vol. 8667. Springer, 194–205.
- [41] Andreagiovanni Reina, Roman Miletitch, Marco Dorigo, and Vito Trianni. 2015. A quantitative micro–macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence* 9, 2–3 (2015), 75–102.
- [42] Andreagiovanni Reina, Gabriele Valentini, Cristian Fernández-Oto, Marco Dorigo, and Vito Trianni. 2015. A design pattern for decentralised decision making. *PLoS One* 10, 10 (2015), 1–18.
- [43] Karina A. Roundtree, Jason R. Cody, Jennifer Leaf, H. Onan Demirel, and Julie A. Adams. 2019. Visualization design for human-collective teams. In *Proceedings of the Human Factors and Ergonomics Society 67th Annual Meeting*. 417–421.
- [44] Michael Rubenstein, Christian Ahler, and Radhika Nagpal. 2012. Kilobot: A low cost scalable robot system for collective behaviors. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 3293–3298. DOI: <https://doi.org/10.1109/ICRA.2012.6224638>
- [45] Erol Şahin. 2005. Swarm robotics: From sources of inspiration to domains of application. In *Swarm Robotics. Lecture Notes in Computer Science*, Vol. 3342. Springer, 10–20.
- [46] Ozgur Koray Sahingoz. 2014. Networking models in flying ad-hoc networks (FANETs): Concepts and challenges. *Journal of Intelligent & Robotic Systems* 74, 1 (April 2014), 513–527.
- [47] Paul Scharre. 2014. Robotics on the battlefield part II: The coming swarm. *Center for a New American Security*. Retrieved January 28, 2021 from <https://www.cnas.org/publications/reports/robotics-on-the-battlefield-part-ii-the-coming-swarm>.
- [48] Alexander Scheidler, Arne Brutschy, Eliseo Ferrante, and Marco Dorigo. 2016. The k -unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics* 46, 5 (2016), 1175–1188.
- [49] Thomas Schmickl and Karl Crailsheim. 2008. Trophallaxis within a robotic swarm: Bio-inspired communication among robots in a swarm. *Autonomous Robots* 25, 1 (2008), 171–188. DOI: <https://doi.org/10.1007/s10514-007-9073-4>
- [50] Thomas Seeley, P. Kirk Visscher, Thomas Schlegl, Patrick M. Hogan, Nigel R. Franks, and James A. R. Marshall. 2012. Stop signals provide cross inhibition in collective decision-making by honey bee swarms. *Science* 335 (Aug. 2012), 108–111.
- [51] Thomas D. Seeley and P. Kirk Visscher. 2004. Quorum sensing during nest-site selection by honeybee swarms. *Behavioral Ecology and Sociobiology* 56, 6 (Oct. 2004), 594–601. DOI: <https://doi.org/10.1007/s00265-004-0814-5>
- [52] Stephen J. Selcon, Ronald M. Taylor, and E. Koritsas. 1991. Workload or situational awareness? TLX vs. SART for aerospace systems design evaluation. In *Proceedings of the Human Factors Society Annual Meeting*, Vol. 35. SAGE Publications, Los Angeles, CA, 62–66.
- [53] Mohamed S. Talamali, James A. R. Marshall, Thomas Bose, and Andreagiovanni Reina. 2019. Improving collective decision accuracy via time-varying cross-inhibition. In *Proceedings of the 2019 IEEE International Conference on Robotics and Automation (ICRA'19)*.
- [54] Gabriele Valentini. 2017. *Achieving Consensus in Robot Swarms: Design and Analysis of Strategies for the Best-of-n Problem*. Springer International.
- [55] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. 2017. The best-of- n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI* 4 (2017), 9. DOI: <https://doi.org/10.3389/frobt.2017.00009>
- [56] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. 2014. Self-organized collective decision making: The weighted voter model. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems*. 45–52.
- [57] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. 2015. Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. 1305–1314.
- [58] Steven G. Vandenberg and Allan R. Kuse. 1978. Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills* 47, 2 (1978), 599–604. DOI: <https://doi.org/10.2466/pms.1978.47.2.599>
- [59] Phillip Walker, Steven Nunnally, Michael Lewis, Nilanjan Chakraborty, and Katia Sycara. 2014. Levels of automation for human influence of robot swarms. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*.

- [60] P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. 2012. Neglect benevolence in human control of swarms in the presence of latency. In *Proceedings of the 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC'12)*. 3009–3014.
- [61] Jan Wessnitzer and Chris Melhuish. 2003. Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In *Advances in Artificial Life*. Lecture Notes in Computer Science, Vol. 2801. Springer, 893–902. http://dx.doi.org/10.1007/978-3-540-39432-7_96
- [62] Robert Wood, Radhika Nagpal, and Gu-Yeon Wei. 2013. Flight of the Robobees. *Scientific American*. Retrieved January 28, 2021 from <https://ssr.seas.harvard.edu/files/sciam2013-robobee.pdf>.

Received April 2020; revised September 2020; accepted December 2020