DISCRETE CONSENSUS DECISIONS IN HUMAN-COLLECTIVE TEAMS

By

Jason R. Cody

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
in
Computer Science

May 11, 2018
Nashville, Tennessee

Approved:
Julie A. Adams, Ph.D.
Yevgeniy Vorobeychik, Ph.D.
Maithilee Kunda, Ph.D.
Jennifer S. Trueblood, Ph.D.
Alexander S. Mentis, Ph.D.

Acknowledgments

I remain indebted to family members, friends, and mentors for the support and guidance I received throughout this research. My wife, Jocelyn, was and is my keystone. She believed in me when I doubted myself and frequently shouldered her career and our children when requirements pulled me away. Her readiness to help me buckle down, or take a break as needed were equally invaluable to this work. I am grateful to my son Logan, who was old enough to know the "Test" (his word) was important, young enough to be hurt anyway, but kind enough to forgive me and let me make it up to him and his younger sister, Caroline. I am grateful to my parents and siblings for their uncanny tendency to reach out when I needed it most.

I wish to thank my advisor, Dr. Julie Adams, for her coaching over the past two and a half years. She is a rare mentor who manages to be simultaneously demanding, responsive, and supportive. I will take her lessons on writing and research to my grave, but only after imparting them to my own students first. I thank my dissertation committee members for the time, guidance, and perspective they provided, which shaped my current research and will influence my future pursuits. I also thank fellow lab members Anton, Jamison, Electa, and Musad whose friendship, conversations, and camaraderie I continue to treasure.

Finally, I am grateful to the United States Army and the United States Military Academy for providing me with this great opportunity. I look forward to returning to my colleagues at the USMA Department of Electrical Engineering and Computer Science.

TABLE OF CONTENTS

List of Tables

List of Figures

List of Symbols

Collective Action Selection

Q_m	A queue of neighboring opinions, $Q_m = \{q_1, .., q_k\}$, where k is the length of the queue.
CAS	A collective action selection problem.
A	The collective of N agents, $A = \{a_1, .., a_N\}$.
Φ	The set of n options within an environment, $\Phi = \{\phi_1, .., \phi_n\}$.
Act	The set of required agent actions.
Act_{assess}	The set of agent assessment actions.
Act_{choose}	The set of agent actions the collective is choosing between.
f_i	A "found" function, where $f_i(\phi_i) = true$ if the $\phi_i \in \Phi^{found}$.
v_i	The value of ϕ_i.
$cost_i$	The cost required to find and assess ϕ_i (e.g., time or resources).
loc_i	The location of $\phi_i \in \mathbb{R}^{dim}$, where dim is the necessary dimension.
h	The option assessment function.
s	An agent's local state.
S	The set of available states.

Collective Model

T	Quorum threshold for sensing a majority of agents in the same state. sort
Ψ	Total population of agents in the collective (Ψ=1).
Ψ_s	Total population of agents in state s ($\Psi_s = A_s/N$).
$\dot{\Psi}_s$	Rate of change in subpopulation Ψ_s.
R_*	Transition rates for transitions $\in \{\gamma_i, \sigma_i, \rho_i, \alpha_i\}$.
r_s^L	Ratio of latent agents in the population in state s.
r_s^I	Ratio of interactive agents in the population in state s.
d_i	The distance between ϕ_i and the decision making hub
d_{max}	The maximum range for the collective.
τ_i	The round trip time for option ϕ_i.
$\bar{\tau}$	The average round trip time.
τ_{max}	The maximum round trip time.
ρ_i	Recruit probability for option i.
γ_i	Discovery probability for option i.
σ_i	Inhibition rate for option i.
α_i	Inhibition rate for option i.
k	The length of the agents' message queues.
ω_s	The quorum sensing rate for state s.
R	The value resolution for the collective best-of-n model.

Individual Behavior Model

P_*	Transition probabilities for transitions $\in \{\gamma_i, \sigma_i, \rho_i, \alpha_i\}$.
U^*	The uncommitted state, where * indicates either the latent (L) or interactive (I) states.
F_i^*	The favoring state supporting ϕ_i, where * indicates either the latent (L) or interactive (I) states.
C_i	The committed state for ϕ_i, in which agents prepare the collective to execute action act_i.
I_i	The initiating state for option ϕ_i.
X_i	The execution state for option ϕ_i.
D_i	The completed state for option ϕ_i.

Chapter I

Introduction

The recent growth in efforts to design, control, and understand large decentralized groups of robots is fueled by the increasing availability, capability, and efficiency of robotic platforms (Kolling et al., 2016). The field of swarm robotics seeks to coordinate large groups of locally sensing and communicating robots in order to produce behaviors that are robust to population attrition, scalable to different population sizes, and flexible to changing environmental conditions (Şahin, 2005). Decision making is a universal capability that is routinely accomplished in social animal groups (Conradt and Roper, 2005) and has been intensely studied in swarm robotics (Brambilla et al., 2013). Common to numerous recent studies is the requirement that robotic collectives solve the best-of-*n* problem, in which the group chooses the best option from a finite set of *n* options (Valentini et al., 2017). Whether the collective chooses between aggregation sites (Arvin et al., 2014; Campo et al., 2011) or future actions (Parker and Zhang, 2009, 2010), its decisions are often influenced by environmental features (e.g., distance) that alter robot interactions and bias the collective towards options that are the easiest to find, evaluate, and reach, but may not be the optimal choice (Valentini, 2017). Collectives that must make decisions, despite biasing environmental factors are likely to be inaccurate and inefficient (Valentini, 2017). This dissertation develops a biologically-inspired collective action selection strategy that enables a robotic collective to choose and implement its best action, despite the presence of environmental bias. The strategy is evaluated for its ability to act independently and as a contributing member of a human-collective team in trials featuring the human supervision of multiple collectives each simultaneously solving a best-of-*n* problem.

The majority of existing collective best-of-*n* strategies either make decisions in specific scenarios that exclude environmental bias (e.g., Hamann et al., 2012; Garnier et al., 2009), or rely on environmental bias in order to discover the shortest paths (Scheidler et al., 2016; Schmickl and Crailsheim, 2008) or the closest sites (Reina et al., 2014). Only two studies (Reina et al., 2015b; Campo et al., 2010) make decisions favoring options based on quality, despite the influence of environmental features (Valentini et al., 2017). Future systems that are unable to overcome bias caused by environmental features will not generalize to the terrain and will require either perfect prior knowledge of the deployment site or direct human supervision and control in order to achieve predictable success. An example of the environmental bias problem is a site selection scenario in which the collective must choose the best site to initiate a collective construction task (e.g., as proposed by Parker and Zhang, 2009). A collective that is unable to compensate for travel cost will be biased to choose whichever acceptable site is closest, even if better sites are within range (as demonstrated by Reina et al. (2015b)). This dissertation proposes the use of cost sensitive control states to enable collectives to compensate for existing bias and to make quality decisions within a maximum exploration range.

Few robotic collective studies combine best-of-*n* decision making with collective behavior changes. Sometimes achieving consensus, or a majority of support, for the best option is sufficient for accomplishing the desired task (e.g., Campo et al., 2011; Schmickl and Crailsheim, 2008). Often, it is more desirable to permit the collective to make a decision and alter its behavior in response to the decision event (e.g., Parker and Zhang, 2010; Wessnitzer and Melhuish, 2003). Transitioning between decision making activities and action implementation is challenging for robotic collectives, due to the limited sensing capability of the collective's individuals and the collective's dynamic communication network. Individuals must estimate when the collective has reached a decision and when it is ready to initiate the chosen task. The quorum sensing process is used by many biological collectives to detect changes in collective state (e.g., Seeley, 2010; Hamar and Dove, 2012; Mallon et al., 2014), which permits individuals to estimate collective state based on recent direct or indirect information exchanges with neighboring individuals. Although biological collectives accomplish action selection when their members are dispersed over a wide area (Mallon et al., 2001; Seeley, 2010), robotic collective studies that enable collective action selection have assumed either a connected network (Wessnitzer and Melhuish, 2003), or a random network enclosed within a small arena (Parker and Zhang, 2010). Robot dispersion is likely to increase population attrition, as robots that are unaware of a decision will be left behind when the collective initiates its chosen action. This dissertation proposes a model that requires sequential quorum sensing to permit individuals to estimate when the collective is ready to transition from decision making to action implementation and from implementation back to a decision making state. The collective is assumed to be dispersed in a wide area and its individuals interact periodically within a shared deliberation area (i.e., a honeybee hive) that is smaller than the collective's total area.

Human-Swarm Interaction (HSI) research has enabled a single human operator to control large numbers of robots in one, or more, artificial swarms by interacting with the group as a single entity (Brown et al., 2016; Kolling et al., 2016). Humans are usually required to make all group behavior decisions for the artificial swarm, which assumes the human maintains global environmental knowledge. Very few studies consider the likely scenario that the human does not have perfect knowledge and must delegate portions of the decision making process to the artificial swarm (Crandall et al., 2017). Humans routinely choose the best of a set of available actions for the swarm to perform (e.g., alter behavior Kolling et al., 2013 or investigate a location Walker et al., 2013). Swarms in current HSI studies; however, are unable to make independent best-of-*n* decisions, even though examples exist in biology (Seeley, 2010; Mallon et al., 2001) and swarm robotics (Valentini, 2017; Reina et al., 2015b). This dissertation evaluates collective action selection in human-collective teams and evaluates the effectiveness of specific control mechanisms and visualization techniques.

I.1 Research Questions

This dissertation focuses on the question: *How can independent robotic collectives and human-collective teams make action selection decisions in dynamic environments?* The solution to this problem requires addressing a number of subordinate research questions, that include:

R_1: How can a collective make efficient and accurate decisions when the problem space is subject to negative environmental bias?

R_2: How can *well-mixing* dynamics be adjusted to improve collective decision accuracy and efficiency?

R_3: How responsive are collective best-of-*n* decision making processes to changes in option values?

R_4: How can a collective transition between decision making and task implementation behaviors?

R_5: What mechanisms permit efficient human control over robotic collective action selection?

R_6: What visualization techniques improve human-collective interaction during collective action selection?

I.2 Outline

Chapter II reviews current research from three relevant areas: collective animal behavior, robotic collective decision making, and Human-Swarm Interaction. Collective animal behavior continues to inspire advances in robotic collective decision making, which in turn offers opportunities for future human-interaction and control protocols. Chapter III presents the formal problem, solution approaches to the first four research questions (R_1 through R_4), and simulation based experimental results. Chapter IV describes solution approaches to the remaining research questions (R_5 and R_6) and the results of a human subjects experiment with an interface evaluating the collective action selection algorithm in human-collective targeting scenarios. Specific visualization techniques and interaction mechanisms are evaluated for their ability to enable human control and understanding of the collective's decision making behavior. Chapter V summarizes this dissertation's contributions and discusses remaining open areas related to this research.

Chapter II

Literature Review

Biological collective behaviors have inspired the development of scalable and robust decision making strategies in decentralized robotic collectives. This chapter reviews biological sources of inspiration, before reviewing the state of the art in robotic collective decision-making and Human-Swarm Interaction.

II.1 Biological Collective Decision Making

Natural phenomena, such as flocking birds and foraging insects, emerge from local interactions of individuals with their environment and neighbors (Bonabeau et al., 1999; Sumpter, 2010). Individuals make decisions that are influenced by stimuli thresholds and regulatory feedback systems, which causes the emergence of intelligent group responses to the environment (Couzin, 2009; Sasaki and Pratt, 2012). Resource scarcity drives these behaviors to balance the energy required to make a decision against the benefit of the outcome.

The majority of collective animal behavior has been categorized into two processes: combined and consensus decision making (Conradt and Roper, 2005). Combined decisions require individuals to determine when their estimate of local stimuli exceeds their response threshold for conducting a specific activity (Wilson, 1984; Bonabeau et al., 1999; Sumpter, 2010). Individuals do not actively coordinate their activities, but coordination emerges from the combination of individual responses to available stimuli (Bonabeau et al., 1999, 1998). Consensus decisions occur when a group must choose a single action from the available actions and incur either a cost in energy to achieve the consensus, or a compromise cost in which individual interests are ceded to those of the group (Conradt and Roper, 2005). Social insect foraging, task allocation, resource distribution and collective construction are examples of combined decisions (Bonabeau et al., 1999). Consensus decisions include flocking, small group hunting and grazing, group navigation, and future nest site selection (Couzin et al., 2005; Seeley et al., 1991; Young et al., 2013).

II.1.1 Combined Decisions

Individual behaviors in combined and consensus decision making processes are the result of nonlinear responses to local stimuli that are affected by regulatory feedback (Sumpter, 2010). Even individuals with high response thresholds, which bias them against a specific task, will perform the task if their estimate of the associated stimuli exceeds their higher response threshold. This phenomenon was demonstrated by Wilson (1984) with *Pheidole pubiventris* colonies, which are divided between workers of significantly different size and task thresholds. The larger ants serve as guards, while the smaller ants perform all other behaviors, including brood rearing, queen care, and foraging. After the smaller ants were removed during an experiment, the larger ants rapidly assumed all colony duties, which indicates the collective's decentralized capacity to maintain a stable state despite changing conditions (Wilson, 1984).

Although combined decisions do not require active behavior coordination, individuals must share information using either indirect or direct interactions (Conradt and Roper, 2005). Individuals indirectly interact by altering the environment in a way that influences the decisions of the next individual, which is a communication technique referred to as *stigmergy* (Bonabeau et al., 1999). Ant corpse management provides an example of stigmergy-based communication (Theraulaz et al., 2003), in which ants tend to place corpses together. When an ant places a corpse in a location, it increases the probability that the next ant will place its load next to the first, which causes corpses to be collected in increasingly larger piles. Pheromone laying is another common indirect communication method and serves as a volatile, global memory in foraging and construction problems (Bonabeau et al., 1999; Sumpter, 2010). Individuals directly communicate through physical contact and deliberate information transfers. Examples of direct communication include the famous honeybee waggle dance used to communicate foraging site locations to other bees (Seeley, 1995) and the tandem runs *Temnothorax rugatulus* ants use to lead foragers to potential sites (Shaffer et al., 2013).

Combined decisions are the least costly of the collective decision making processes and often solve problems of distribution, allocation, regulation, and synchronization (Sumpter, 2010). These decision making behaviors permit the collective to reach and maintain stable states, which are necessary for collective survival in dynamic environments (Şahin, 2005). Combined decisions are important, but many problems facing social animal groups require a single unified response.

II.1.2 Consensus Decisions

Consensus decision making requires individuals to actively coordinate their behaviors in order to enable the collective to converge to a single action (Conradt and Roper, 2005). Bird flocking and fish schooling are classic examples and emerge from individuals' effort to balance personal requirements (e.g., nourishment), with the security of remaining in a cohesive group (Sumpter, 2010; Couzin et al., 2002; Attanasi et al., 2014). Shared consensus decisions are those made by several members of the group, while unshared decisions are made by a dominant individual (Conradt and Roper, 2009). The latter are efficient for small group decisions and individuals assume a leadership role as a result of their assertiveness and location within the existing dominance hierarchy (Conradt and Roper, 2009). This chapter focuses on shared consensus decisions, which enable large groups to make decisions without relying on specific individuals. Both continuous and discrete consensus decisions are common in animal groups (Valentini, 2017). Continuous consensus decisions require the group to achieve consensus for one of an infinite number of options, while the discrete consensus decisions require the group to choose from a finite set of options. Collective solutions of the best-of-*n* problem are a special type of discrete consensus achievement that ensures the option chosen is also of the highest value to the group.

II.1.2.1 Continuous Consensus Achievement in Animal Groups

A transformational model of bird flocking behavior, developed by Reynolds (1987) and inspired by existing fish schooling research (Aoki, 1982) demonstrated that simple repulsion and attraction rules enabled the birds to accomplish complex coordinated movements. This model was extended by Couzin et al. (2002) to incorporate zones of repulsion, orientation, and attraction and accurately recreated distinct collective movement states observed in fish schools. Individual fish respond nonlinearly to changes in their neighbor's behaviors (Couzin et al., 2005). When the number of peers altering their behavior exceeded a predetermined threshold, the probability of an individual adopting the same behavior greatly increased, similar to the individual stimuli responses described in Chapter II.1.1. This response mechanism effectively reduces the impact of individual error on collective consensus decisions and increases the probability that a group will make decisions in favor of the majority opinion (Ward et al., 2008).

Several communication models for continuous consensus achievement have been proposed for swarm movement. The metric model places an individual's neighbors in repulsion, orientation, and attraction zones (e.g., Couzin et al. 2002; Reynolds, 1987; Aoki, 1982) in which agents respond to neighbors in these zones defined by straight line distances. The topological model restricts agent interactions to a fixed number of neighbors, regardless of distance between neighbors and has been demonstrated to improve accuracy over the metric model in simulating biological swarm movement (Ballerini et al., 2008; Abaid et al., 2015). Another model, developed by Ginelli and Chaté (2010), restricted agent alignment to neighbors within the agent's current region of a Voronoi diagram. A visual model that uses agent sensory information to define agent neighbors was shown to outperform the other three models in accounting for fish movement. (Strandburg-Peshkin et al., 2013). The visual model was extended in order to account for oscillations in relative speed between fish that further improves movement coordination (Swain et al., 2015). Although simulated and experimental data with biological agents increasingly supports the visual model, communication methods for artificial swarms remains an area of open research (Strandburg-Peshkin et al., 2013; Haque et al., 2016).

II.1.2.2 Discrete Consensus Achievement in Animal Groups

Animal groups often choose between a finite set of available options. When environmental features restrict a flocking group to a limited number of possible headings, the decision strategy described in Chapter II.1.2.1 supports a discrete choice. A model of group decisions during animal movement demonstrated the effect of different individual heading preferences within a group (Couzin et al., 2005). When heading preferences are close, the group effectively averages them, but when the preferred headings are different, the group commits to the option with the greater support. This decision capability was analyzed by Hartnett *et al.* (2016) using a physical spin system model. The model favored majority opinions; however, the number of uninformed individuals and the degree of collective information pooling increased the probability that a minority group with a strong opinion altered the collective's decision.

Information pooling is a primary component of biological solutions to the collective best-of-*n* problem, in which the benefit of the decision to the entire group is incorporated into the decision process and a *consensus choice* is made (Golman et al., 2015; Valentini, 2017). These decision making strategies require a significant investment of energy, which makes them rare in nature, as such, these types of strategies have only been studied extensively in two species (Franks et al., 2002; Sasaki and Pratt, 2012; Seeley, 2010). The *Temnothorax albipennis* ant and the *Apis mellifera* honeybee solve the best-of-*n* problem when they emigrate to a new location. Potential future sites are evaluated by scouting members of each collective. Site quality determines the strength of scout recruitment efforts and often the likelihood of the site's selection by the group (Franks et al., 2002). The consequence of selecting the wrong site is a severe drop in the collective's chances of survival, which has resulted in a decision making process with negligible conflict of interest between individuals (Seeley, 2010). Despite the insects' aligned intent, they face significant challenges. First, individual insects are unable to do better than guess randomly after comparing multiple sites (Sasaki and Pratt, 2012), which means collective decisions do not rely on direct site comparisons by individuals. Second, individuals are unable to observe the collective's global state during the decision making process. Individuals estimate the collective state, based on local observations, but the decision process must prevent incorrect estimates from causing the group to emigrate to a site prematurely.

II.1.2.3 *Apis mellifera* Honeybee Nest Site Selection

The *Apis mellifera* nest site selection process, as detailed by Seeley and collaborators (1978, 2010, 2012), continues to influence numerous fields. Future nest sites must be selected to optimize honey production, which guarantees colony survival throughout the winter (Seeley, 2010). The ideal nest sites are 6.5 meters above the ground with a cavity volume of 40 liters and have a southern facing 15cm^2 entrance near their base (Seeley and Morse, 1978). The honeybees identify dozens of alternative nest sites in a wide geographic area (150km^2), evaluate each according to these metrics, and choose the best site within a limited amount of time.

Every spring, healthy honeybee colonies emit at least one swarm of roughly 15,000 workers with their queen (Seeley, 2010). The swarm converges temporarily at a location near the original hive, while 3-5 percent of its members, the honeybee scouts, explore for potential nest sites. Scouts that discover acceptable sites return to the swarm and recruit other scouts using a “waggle dance” to communicate their site’s location (Seeley, 2010). The number and duration of the scout’s dances depends upon the quality of its site and affects its probability of recruiting other bees (Seeley, 2010). Recruited scouts travel to the site and evaluate it in order to determine their own recruiting strength. The scouts also inhibit other scouts that are recruiting for different sites. Inhibition signals, called *stop* signals, cause the recipient to eventually cease recruitment and inhibition. Inhibition accelerates the dismissal of poor quality sites and shortens the debate length (Seeley et al., 2012). During the scouts’ deliberation, the remainder of the swarm ensures the queen’s survival and awaits a decision.

Scouts that detect the presence of twenty to thirty peers investigating their site estimate that the swarm has reached a decision and alters their behavior to prepare the swarm to emigrate (Seeley, 2010). A quorum of this size is highly unlikely unless more than half of all scouts have assessed and support the site. Honeybees have optimized their decision making process in order to minimize the time spent as an unprotected swarm, while maximizing the likelihood of finding the best possible site. Quorum sensing enables the collective to balance the speed and accuracy of this decision making process (Franks et al., 2002; Seeley, 2010).

High quality nest sites induce higher probability of recruiting and more inhibition of competing sites, which increases support for high quality sites until one site reaches a quorum and scouts begin signaling for movement (Seeley, 2010; Seeley et al., 2012). Analysis of the honeybee nest site selection process has revealed it to be value-sensitive, which means that the ability to distinguish between sites depends upon the difference in site values and the mean site value (Pais et al., 2013). Value-sensitivity also enables the collective to delay committing to poor sites, which allows them to quickly converge to a higher quality site if one is discovered early enough in the decision process (Passino, 2010; Pais et al., 2013).

II.1.2.4 *Temnothorax albipennis* Ant Nest Site Selection

The *Temnothorax albipennis* ant conducts an emigration process that is executed without colony pheromone trails, reliably selects the best of available future nest sites, and depends upon quorum sensing, nest site quality estimation, and feedback systems (Mallon et al., 2001). The ants form colonies in tiny rock crevasses and frequently emigrate when the current nest site becomes uninhabitable. The scouting population, which makes up a full third of the colony, assesses potential nest sites for capacity, cavity height, entrance size, and interior light level (Mallon et al., 2001; Sasaki and Pratt, 2013). The ants face a similar problem to the honeybees and must choose the best site from several within a search space based on the described metrics.

Ant colony scouts begin exploring for a new site when the current site becomes overcrowded or is damaged (e.g., exposed due to environmental change). Upon locating an acceptable site, a scout recruits other scouts after a delay period that is inversely proportional to the site’s quality. The delay period prevents mediocre sites from being chosen before other sites have been considered (Pratt et al., 2002). Scouts recruit one peer at a time by performing a tandem run, in which the informed scout leads its recruit to the discovered site, at which the recruit performs an individual assessment (Mallon et al., 2001). Tandem runs continue until a scout detects a quorum of peers at the site. Upon sensing a quorum, a scouting ant is believed to identify the new location as its home and immediately changes its recruiting strategy from tandem runs to transports (Pratt et al., 2002). Scouts physically pick up other ants and carry them to the new site during transports, which is significantly faster than tandem runs (Mallon et al., 2001). As a collective, the ants deliberate and emigrate simultaneously (Franks et al., 2002). Although the presence of transports indicate a colony decision, multiple quorum detections occasionally result in transports to multiple alternative sites (Mallon et al., 2001).

Multiple simultaneous quorums are not disastrous to the *T. albipennis* ants. The *Temnothorax curvispinosus* ants, which have an identical nest site selection process to *T. albipennis* ants, adjust the required quorum number, depending on the circumstances of the colony's emigration (Pratt and Sumpter, 2006). If the colony is exposed due to the sudden loss of its protective roof, the ants are more likely to sense multiple quorums and move the original colony to many sites, only to retrieve all distributed ants and bring them to the chosen nest site later. Emigrations initiated due to overcrowding at the original site are less likely to result in dispersal to multiple new sites. Multiple quorums potentially enhance the survival of the colony when emigration occurs, while the original colony is vulnerable (Pratt and Sumpter, 2006).

II.1.2.5 Nest Site Selection Process Comparison

The collective best-of-*n* decision making processes of *Temnothorax albipennis* and *Apis mellifera* have been tailored by evolution to efficiently solve the nest site selection problem given the capabilities, limitations, life cycle, and environment of each species. Six common components of these decision strategies exist (Franks et al., 2002): information gathering, evaluation, deliberation, consensus building, choice and implementation.

Both species delegate the decision making processes to a subset of dedicated individuals that explore for future sites and select the best one available. Site quality positively influences recruitment, which accelerates the growth of support for optimal sites over mediocre sites. When only mediocre sites are discovered, one of them is eventually chosen, unless an optimal site is discovered early enough in the decision making process (Pais et al., 2013). Quorum sensing enables both species to recognize when a decision has been reached and initiate emigration behaviors (Franks et al., 2002; Sumpter and Pratt, 2009). Site comparisons occur as a result of the relative number of supporting agents, rather than requiring individuals to compare available value estimates (Seeley and Morse, 1978; Mallon et al., 2001).

Several important differences exist between theses emigration strategies (Franks et al., 2002). Honeybees create expensive and immobile honeycomb within their colonies and are unable to recover from multiple simultaneous migrations, like the *T. albipennis* ants. Honeybees search a large area for nest sites, whereas *T. albipennis* ants may find several sites within a meter of their original nest (Mallon et al., 2001; Seeley, 2010). The process duration also differs between the species. The honeybee process lasts two to four days, while the *T. albipennis* ants complete their emigration within a few hours. Additionally, the *T. albipennis* ants implement their emigration concurrently with their deliberation, while the honeybees initiate emigration only after a decision is reached (Seeley, 1982; Franks et al., 2002). Finally, site quality affects scouts from each species differently. *T. albipennis* ant scouts delay initial recruiting based on site quality, but do not stop recruiting after the delay, although they may switch site allegiance (Franks et al., 2002; Pratt et al., 2002). Site quality does not induce a delay in the honeybees, but does influence the vigor of the honeybee scouts' recruiting and inhibition efforts (Seeley, 2010; Seeley et al., 2012).

Early efforts to develop similar decision making processes in robotic collectives have identified challenges associated with interaction rates and distances between target sites (Reina et al., 2015a). All steps of the nest site selection process are tuned to each other by evolution. The recruiting behaviors and quorum sensing strategies implicitly depend on scout movement and site discovery rates and have been developed through a long evolutionary process. Generalizing these techniques requires discovering the underlying relationships between these behaviors and the environment in order to decouple the decision making process from any one specific environment. This challenge is an example of the "inverse problem" (Berman et al., 2007), which is defined as the difficulty of developing individual behaviors that produce desired collective results.

II.2 Robotic Collective Decision Making

The decision making processes described in Chapter II.1 have inspired significant research into the development of similar behaviors in robotic swarms (Brambilla et al., 2013; Kolling et al., 2016; Valentini et al., 2017). Robotic swarms are large groups (> 50) of robots with limited sensing, communication, and computation capabilities whose interactions cause the emergence of a group behavior (Şahin, 2005; Hayes and Adams, 2014). Recent robotic swarms demonstrate emergent behaviors that scale well with changes in group population, react flexibly to dynamic environments, and are robust to the effects of robot attrition (Brambilla et al., 2013). These characteristics make swarm robotics a reasonable approach to solving future problems in numerous domains including: medical (Mermoud et al., 2010), military (Wessnitzer and Melhuish, 2003; Coppin and Legras, 2012), transportation (Campo et al., 2010), and space exploration (Kolling et al., 2016; Truszkowski et al., 2004). A recent survey segregates swarm robotics into three primary behaviors (Brambilla et al., 2013), which include decision making (e.g., combined decisions (Krieger and Billeter, 2000) and consensus decisions (Montes de Oca et al., 2011)), spatial organization (e.g., pattern formation (Spears et al., 2004) and aggregation (Campo et al., 2011)), or navigation (e.g., flocking (Reynolds, 1987)).

This dissertation draws an important distinction, in terms of human-collective interaction, between swarms and collectives. A swarm is a decentralized group whose individual decisions are based on distributed individual preference and local information, which is common in coordinated movement in biological groups (Couzin et al., 2002). Swarm behaviors include consensus decision strategies (see Chapter II.1.2) that result in the group's selection of one option from either an infinite (e.g., continuous consensus achievement) or a finite (e.g., discrete consensus achievement) set. A *collective* consists of individuals whose decisions are biased to favor the needs of the group, which is facilitated by information pooling within a shared space, or decision-making hub (Valentini, 2017). Collective behaviors include combined decision (see Chapter II.1.1) and best-of-*n* decisions similar to the emigration behaviors described in Chapters II.1.2.3 through II.1.2.5.

The development of best-of-*n* strategies in robotic collectives is an area of intense recent study (Wessnitzer and Melhuish, 2003; Parker and Zhang, 2009; Scheidler et al., 2016; Valentini, 2017; Reina et al., 2015b). These strategies are analogous to individual decision making and action selection (Barron et al., 2015; Parasuraman et al., 2000) when used to choose the best action for execution. Collective decision making is facilitated by a combination of information pooling within the collective's shared space, or *decision making hub*, and robot behavior modulation that drives the system to favor one action (Valentini, 2017). Collective action selection is either implicit or explicit, as shown in Figure II.1. Implicit action selection accomplishes the chosen action during the decision process (Campo et al., 2011; Montes de Oca et al., 2011; Scheidler et al., 2016), while explicit action selection includes decision detection in the consensus achievement process in order to enable the collective to execute a chosen action after the decision (Wessnitzer and Melhuish, 2003; Parker and Zhang, 2009, 2010). Explicit action selection enables robotic collectives to choose future actions and break up complex missions into sequences of decisions and tasks (Parker and Zhang, 2010).

Concurrently with the research into the design, performance, and implementation of robotic collectives (Brambilla et al., 2013), the field of Human-Swarm Interaction has emerged to examine the challenges to human control and interaction that impede practical robotic swarm deployment. The Human-Swarm Interaction research has only recently begun to examine human interaction with robotic collectives capable of implicit and explicit action selection (Coppin and Legras, 2012; Crandall et al., 2017). The human's role, in most human-swarm interactions, has focused on influencing the overall swarm's behavior, often via influencing individual swarm members (Kolling et al., 2016). Recent research suggests that human-collective teams alter this relationship by permitting the human to share control with the collective during decision making, which is particularly useful when neither the human, nor the collective, has access to global information regarding the environment (Crandall et al., 2017). Collective action selection strategies offer two advantages over typical human-swarm interaction. First, collectives use decision strategies that can augment human-collective decision making, or be used to make decisions in the absence of human influence. Second, robots in the

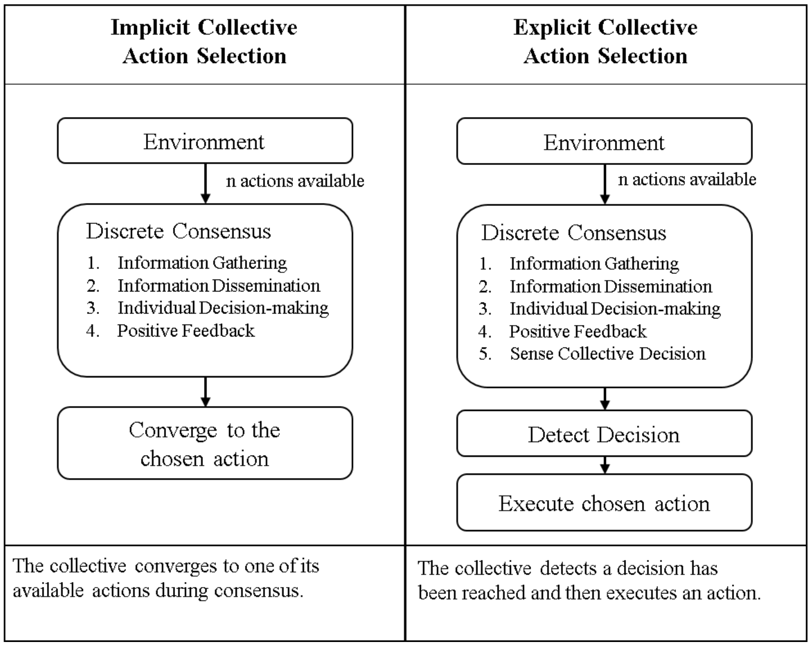

Figure II.1: Implicit collective action selection occurs as a collective converges to one of a set of activities, while explicit collective action selection requires an activity to be chosen prior to execution.

collectives attempt to uniformly share information, which extends the reach and duration of human influence. Collectives also introduce new challenges related to their complex decision dynamics (Valentini, 2017; Reina et al., 2015b) and their increased information requirements necessary to achieve transparency (Chen et al., 2017).

II.2.1 Implicit Action Selection in Robotic Collectives

Implicit action selection is accomplished using collective best-of-*n* strategies in which the collective is required to converge to one of its current actions. Common approaches to the development of these systems include ad hoc (e.g., Campo et al., 2011; Campo et al., 2010) and opinion-based (e.g., Montes de Oca et al., 2011; Reina et al., 2014; Reina et al., 2015b; Valentini et al., 2014; Valentini et al., 2015) approaches. The former often define simple behaviors that cause the collective to converge to a specific location (Campo et al., 2011; Hamann et al., 2012; Arvin et al., 2014) or path within the environment (Schmickl and Crailsheim, 2008). Ad hoc approaches accomplish implicit action selection, since convergence results in the execution of the desired action. Opinion-based solutions are more general and require robots to disseminate internal opinions corresponding to options available to the collective (Valentini, 2017). This chapter focuses on opinion-based best-of-*n* strategies, since they apply to numerous behaviors and uniquely enable both implicit and explicit collective action selection.

A general opinion-based collective best-of-*n* model, developed by Valentini (2017), consists of three components: exploration and dissemination control states, individual decision making mechanisms, and modulation of positive feedback. Exploring robots identify and evaluate the available options, while disseminating robots share opinions with neighboring robots within the disseminating state. Disseminating robots communicate opinions to other disseminating robots within a decision making hub. Restricting communication to a decision making hub eases the challenges of ensuring that collective opinions are *well-mixed*, meaning the probability of any two disseminating robots interacting is approximately equal (Kampen, 2007; Valentini, 2017). Individual decision making mechanisms determine how robots alter their opinions. Finally, positive feedback modulation biases the deliberation between options towards the best available option (Valentini, 2017). Positive feedback modulation is either direct or indirect. Direct feedback modulation requires individual robots to change behavior, based on option value in order to bias the collective towards the best option. Indirect feedback modulation relies on environmental features to bias the decision process (Valentini, 2017).

The general model (Valentini, 2017) easily describes previous successful approaches to implicit action selection (Montes de Oca et al., 2011; Scheidler et al., 2016; Valentini et al., 2014, 2015). Indirect modulation in collectives has been used to generate convergence to the shortest available path (Montes de Oca et al., 2011; Scheidler et al., 2016). Selective foraging problems, in which the collective must choose the best available site, have been solved using both indirect and direct modulation (Valentini et al., 2014, 2015; Reina et al., 2014). A comparison of collective best-of-*n* strategies, conducted by Valentini, Hamann, and Dorigo (2014; 2015), demonstrated that the selected decision strategy had a significant effect on convergence time (Valentini et al., 2015). The majority rule model, in which agents adopt the most supported opinion of their neighbors, was shown to be faster, but less accurate than the voter model, in which robots adopt the opinion of a random neighbor (Valentini et al., 2015). Further, increasing the size of the robot's neighborhood, by increasing its communication range, reduced decision time and accuracy (Valentini et al., 2015). While the influence of option cost (e.g., distance) and value on collective decision making strategies have been investigated, very little research has examined scenarios in which the environmental cost of preferable options is greater than the cost of those that are mediocre (Valentini, 2017).

A similar modular design for collective best-of-*n* decisions was developed by Reina et al. (2015b) and adopted directly from the emigration strategy of the *Apis mellifera* honeybee (see Chapter II.1.2.3). The honeybee nest site selection process was chosen, because it incorporates discovery of new options during deliberation and is value-sensitive (Reina et al., 2015b). Similar to Valentini's modular approach, Reina et al.'s (2015b) design restricts information dissemination to a decision making hub. The model includes the four state transitions exhibited during honeybee nest site selection: discovery, abandonment, recruitment, and inhibition (Reina et al., 2015b). Support for collective options increase as options are discovered by robots and as robots are recruited to evaluate the available options. Site support decreases as robots favoring a site spontaneously abandon it, or when a robot favoring a site is inhibited during dissemination and transitions to the opinion-neutral state. While implementations of Valentini's model have achieved a stable equilibrium at consensus for available options (Valentini et al., 2014, 2015; Valentini, 2017), Reina et al.'s (2015b) design achieves equilibrium when the discovery and recruitment rates are balanced with the abandonment and inhibition rates, which results in a stable majority favoring the highest valued option (Pais et al., 2013).

Both design strategies are the first modular approaches to opinion-based collective best-of-*n* decision making that are generally applicable to numerous scenarios. Several open questions remain that directly affect future use of these decision making processes for human-collective teams. First, three factors bias collective decision making accuracy and decision time: option discovery, option cost, and option value (Passino, 2010; Reina et al., 2015b). These factors refer to the timing of an option's identification by the collective, the energy, or resources required to evaluate the option, and the evaluated quality of the option, respectively (Valentini, 2017; Reina et al., 2015b). Depending upon the system's goals, one, or more of these factors is likely to be preferable to the other two. Collective decision making processes are biased in favor of the options that are discovered early or have a low cost, even when the system's intent is to favor options of the highest quality (Reina et al., 2015b; Cody and Adams, 2017b; Crandall et al., 2017; Valentini et al., 2017). This dissertation proposes the term *disparity resilience* to describe a collective's indifference to one biasing factor when favoring another. Cost disparity resilience, in a scenario requiring the collective to choose between options based on value alone, is the difference between option costs that the collective will effectively ignore when making a decision based on option value. Human operators interacting with collective best-of-*n* strategies will need to understand the collective's disparity resilience in order to predict the conditions that will require human intervention in order to avoid inaccurate decisions due to one or more undesirable biases. Only a few recent studies demonstrate methods of improving collective disparity resilience (Reina et al., 2015b; Cody and Adams, 2017b; Valentini, 2017). Reducing the frequency of robot opinion exchanges improved, but did not eliminate the cost and discovery disparity resilience (Reina et al., 2015b). Another proposed approach penalizes rapidly discovered options with dissemination delays (Cody and Adams, 2017b) and nearby options with lower dissemination frequencies (Cody and Adams, 2017b; Valentini, 2017). Second, option biasing factors are expected to change during the decision process in realistic problem domains. This problem was explored in ad hoc approaches (Schmickl and Crailsheim, 2008; Campo et al., 2011) and for task completion decisions (Parker and Zhang, 2010), but not within the context of general best-of-*n* problems. Collective vigilance toward changes in an option's desired metric (e.g., value) establishes the collective's responsiveness to the environment and to changes introduced by a human, but has not been fully examined. Finally, the chosen decision strategy used by individual members of the collective affects decision time and decision accuracy (Valentini et al., 2015). The implications of accuracy and timing effects of decision making strategies suggest a trade-off between the chosen decision making strategy and a human's workload. The human is likely to increasingly intervene as collective performance decreases.

II.2.2 Explicit Action Selection in Robotic Collectives

Few robotic studies have explored explicit action selection (Wessnitzer and Melhuish, 2003; Parker and Zhang, 2009, 2011), which combines collective best-of-*n* strategies with decision detection mechanisms similar to the emigration behaviors described in Chapters II.1.2.3 through II.1.2.5. Collectives that make future behavior decisions are able to synchronize the completion of a sequence of distinct tasks (Parker and Zhang, 2010) or alter behavior in response to a drastic environmental change (Wessnitzer and Melhuish, 2003). Explicit action selection enables collective decisions related to future actions and collective task sequencing. The primary components of an explicit collective action selection strategy include: 1) discrete consensus achievement strategy, 2) quorum detection, and 3) transition from deliberation to execution.

The first robotic collective capable of explicit action selection required a predatory robotic collective to hunt and disable two different prey robots in sequence (Wessnitzer and Melhuish, 2003). The predatory collective used a majority-rule consensus process to target one of the prey robots. Once the collective surrounded the target, the target was artificially immobilized. Robots detecting that the prey had stopped moving emitted a hormone-inspired message that propagated through the swarm, increasing in strength when received by robots that also detected the immobilized prey. The message contained a timer that enabled the collective's synchronized target switching. The algorithm's effectiveness was demonstrated in simulations with robot collectives with up to 50 robots and hardware experiments with five robots. Although explicit action selection was achieved, the predatory collective randomly selected its target without considering either an intrinsic target value or environmental cost.

The first explicit action selection experiment to consider intrinsic option values was directly inspired by the *Apis mellifera* honeybee (Seeley, 1982; Seeley and Buhrman, 2001) and *Temnothorax albipennis* ant (Mallon et al., 2001) emigration strategies (Parker and Zhang, 2009). During the consensus achievement process, robots used a queue of voting messages to monitor the number of recently encountered robots that shared their opinion. Robots detected a quorum when the ratio of agreeing votes was greater than or equal to the given threshold, called the *quorum threshold*. Robots that sensed a quorum immediately transitioned to a committed state and began transmitting commit messages. Recipients of commit messages acknowledged the message when the recipient joined the committed population, otherwise the robot was silent. Finally, committed robots transitioned to a finished state when they failed to receive acknowledgments after a predefined timer expired. Simulations and experiments with up to 15 robots achieved 80% accuracy in choosing the better of the two available options and reached a consensus in every trial. Increasing the quorum threshold decreased the probability of a falsely detected quorum, increased decision accuracy, and increased decision time.

One challenge that the approach faced was large convergence times when the collective had to choose between options of equal value (Parker and Zhang, 2011). Choosing between equally valued options is a common challenge and a critical ability for collectives, which is called *symmetry breaking* (Hamann et al., 2012). Without centralized coordination, a collective is unable to simply make a random selection between available options. Biological and artificial collectives break symmetry when their positive feedback mechanisms amplify errors in option assessments and interactions, which drives a random decision between equally valued options (Hamann et al., 2012). A *frustration* mechanism was introduced that increased a robot's recruitment probability every time the robot failed to recruit a neighbor, due to the neighbor already sharing the first robot's opinion (Parker and Zhang, 2011). Analysis of the frustrated model revealed that this mechanism was sufficient to enable the collective to break symmetry when necessary.

The previously described decision making strategy was also used as a mechanism for solving the *unary* decision problem, in which a collective executing a task decides when the current task is complete and transitions to the next task (Parker and Zhang, 2010). The collective used quorum detection to decide between continuing the current task or initiating the next task (Parker and Zhang, 2010). Quorum detection noise was identified as an important consideration in determining an appropriate quorum threshold for the collective (Parker and Zhang, 2009, 2010). Duplicate quorum detection interactions and poor population mixing contribute to quorum detection noise (Parker and Zhang, 2010). The proposed solutions to the problem included selecting quorum thresholds and queue lengths sufficient to account for duplicated information, or delay agent interactions in order to reduce the likelihood of quorum detection noise (Parker and Zhang, 2010). Methods of determining the probability of quorum detection noise, or the necessary delay period were not defined, or evaluated.

These initial investigations of collective explicit action selection offer numerous directions for future research, with and without human interaction. First, it is unclear what mechanisms ensure quorum accuracy, which compares the quorum threshold to the actual percentage of the population supporting a chosen action when a decision is detected. False quorum detections, which result from poor quorum accuracy, reduce decision accuracy (Parker and Zhang, 2009) and risk fragmenting the collective when simultaneous quorum detections occur for competing options (Cody and Adams, 2017a). Although false quorum detection probability is low (Parker and Zhang, 2009), the probability of such an outcome increases with population size and robot interaction frequency (Cody and Adams, 2017a; Parker and Zhang, 2010). Both sources of biological inspiration for Parker and Zhang's strategy occasionally fragment, which results in either a costly recovery process (Pratt and Sumpter, 2006) or irreparable dispersion (Seeley, 2010). Mechanisms to resist and recover

Topic	**Research Question**
Implicit Action Selection	How can disparity resilience be improved in collective best-of-*n* decisions?
	How responsive are implicit action selection strategies to changes in option quality?
	How do different individual decision techniques affect implicit action selection performance?
	How does environmental bias influence the maximum number of options in collective action selection problem?
Explicit Action Selection	How can the probability of collective fragmentation be reduced?
	What mechanisms enable collectives to recover from fragmentation?
	What effect do quorum detection noise reduction strategies have on collective decision speed?
	How effective is explicit collective action selection when the collective must move between decisions?

Table II.1: Open research questions for robotic collective action selection strategies.

from collective fragmentation remain an open area of research. Poor quorum accuracy and fragmenting are critical concerns to future human-collective interactions. False quorum sensing reduces the ability of the human to predict when the collective will alter behavior or make a decision. Inadvertent collective fragmentation will require the human to manage different groups, which will increase the human's workload and potentially detract from the human-collective's original task.

A similar problem arises when using explicit action selection in task sequencing (Parker and Zhang, 2011). Neglect benevolence is the concept that swarm performance often improves if human input is delayed (Walker et al., 2012a), which also applies to explicit action selection. This dissertation proposes that a trade-off exists between the speed of a collective's transition from action deliberation to action execution and the likelihood that the collective fragments. Collective fragmentation is more likely when the collective rapidly transitions from deliberation to execution prior to all robots being informed of the collective's decision. Future human-collective teams must respect the collective's instability during this transition and delay human input in order to avoid losing robots or fragmenting the collective prior to execution.

II.2.3 Summary

Robotic collective decision making enables large decentralized groups of relatively simple systems to respond to changes within the collective's environment. These decision behaviors, which include combined (e.g., task allocation) and either continuous (e.g., flocking) or discrete (e.g., nest site selection) decision making, remain active research areas in the field of robotic collectives. This review examined the state of the art for implicit and explicit action selection in robotic collectives. Several open research questions for each of these behaviors are summarized in Table II.1.

Most implicit action selection strategies focus on decisions between a limited number of statically valued options, without any consideration of negative environmental bias (Valentini et al., 2017). Understanding disparity resilience in implicit action selection strategies will enable future collectives to prioritize and reduce the relative influence of biasing environmental factors. This capability is particularly important when the collective is subject to human supervision and its priorities must be adjusted based on changing requirements. Opinion-based models have not been evaluated under conditions of dynamically changing option qualities. The vigilance of implicit action selection strategies to changing option qualities is important to understand, since robotic collectives must be able to operate under changing conditions. Different individual decision making techniques (e.g., majority rule (Castellano et al., 2009)) affect decision speed and accuracy (Valentini et al., 2015), but have not been compared in scenarios for which options are discovered during decision making. Finally, implicit action selection strategies have only been evaluated with a small number of available options (Valentini et al., 2017). Since collectives have a finite population size, there is an expected limit to the number of options the collective is able to reasonably consider at a given time. Some research into this problem has been conducted (Reina et al., 2017), but this limit is expected to also be affected by the existing environmental bias. The number of options a collective is capable of handling will likely change depending upon the actual cost of the available options.

Research related to explicit collective action selection remains limited, but current efforts have revealed numerous areas that require further investigation. Quorum thresholds and quorum detecting mechanisms control the speed and accuracy of decisions; however, false quorum detection has the potential to fragment the collective. Collective fragmentation has been observed in biological collectives (Seeley, 2010; Pratt and Sumpter, 2006), but has not been evaluated in robotic collectives. The possibility of collective fragmentation limits the use of the quorum threshold to control decision speed and accuracy, since lowering the quorum threshold is expected to increase the probability of collective fragmentation. Future investigation of mechanisms to recover from and reduce the likelihood of collective fragmentation is necessary to mitigate this risk. Reducing quorum detection noise is another open research area. Increasing the period between robot interactions (Parker and Zhang, 2010), or using agent identification to discard duplicate messages (Hamar and Dove, 2012) have each been proposed, but not evaluated in explicit action selection research. Finally, existing explicit collective action selection strategies have not been evaluated in a series of decisions, or in situations that require the collective to change its location (Parker and Zhang, 2010). Collectives that change location between decisions are more likely to fragment by leaving portions of the collective behind. Design and evaluation of systems capable of sequential decisions will establish whether these systems can be expected to operate for long periods of time and in large areas.

II.3 Human Interaction with Collective Decision-Making and Action Selection

Defining human-collective interaction requires exploring collective autonomy, transparency, and control. Two existing studies consider human-collective interaction in task allocation (Coppin and Legras, 2012) and site selection (Crandall et al., 2017), but several of the strategies described in Chapters II.2.1 and II.2.2 have yet to be considered for human interaction with a collective. Considerably more work has been conducted in swarm autonomy (Walker et al., 2013; Coppin and Legras, 2012), transparency (Walker et al., 2016; Adams et al., 2018) and control (Kolling et al., 2016), which serves as a starting point for identifying the unique challenges of human-collective interaction. Collectives make group decisions about current (see Chapter II.2.1) and future (see Chapter II.2.2) behaviors based on often unobservable and complex collective decision making dynamics. The advantage of collective decision making processes is that when humans and collectives have similar assessment requirements for the available options, the collective is capable of recommending solutions to the human or making independent decisions. The complex dynamics of collective action selection strategies introduce the need to provide human operators with transparency into the collective reasoning (Adams et al., 2018; Chen et al., 2017) and control mechanisms in order to interact with both the physical (Kolling et al., 2016) and decision making (Crandall et al., 2017) behaviors.

II.3.1 Collective Autonomy

Relatively little research has investigated levels of autonomy in human interaction with collectives (Adams et al., 2018; Crandall et al., 2017; Coppin and Legras, 2012). The levels of autonomy (Sheridan and Verplank, 1978) were intended to describe the independence of a computer teamed with a human, which range from the lowest level in which the human makes all decisions, to the highest level in which the computer makes all decisions. Other levels of autonomy have been proposed (e.g., Proud et al., 2003) and researchers have since examined different autonomy levels for specific decision making functions (e.g., Parasuraman et al., 2000; Endsley and Kaber, 1999) as well as methods of adaptable autonomy (e.g., Coppin and Legras, 2012). The following open research questions are considered:

- What levels of autonomy are available for human interaction with collective action selection?
- What collective action selection features enable low, intermediate, and high levels of autonomy for adaptive human-collective interaction?
- What evaluation criteria assist designers in determining an appropriate level of autonomy for different collective action selection strategies?

The initial focus is on the first two questions by identifying the possible levels of autonomy for collective action selection and comparing these to the limited existing human-collective interaction research. The last question is answered by examining the known limitations of collective action selection strategies described in Chapters II.2.1 and II.2.2 in order to recommend evaluation criteria to enable humans to determine appropriate levels of autonomy.

II.3.1.1 Existing Research in Swarm and Collective Autonomy

Human operators in existing human-swarm interaction research often rely on highly independent swarm behaviors in order to search (Walker et al., 2012a; Goodrich et al., 2011; Walker et al., 2013), cover (Kira and Potter, 2009; Jung and Goodrich, 2013; Jung et al., 2013), or traverse (Nagavalli et al., 2017) an area. These analyses require the human to interpret information from the swarm and make decisions about the swarm's current and future behavior, including which target to choose (Goodrich et al., 2011), formation to adopt (Kira and Potter, 2009), or path to take (Nagavalli et al., 2017). Some human-swarm interaction research has begun to examine different levels of autonomy in human-swarm teams (Kolling et al., 2016). One proposed approach permitted the human to switch between high and low levels of robot autonomy in order to accomplish a target search (Walker et al., 2013). At high autonomy, robots independently approached targets and avoided obstacles or other robots. At low autonomy, humans selected robots in the swarm and directed them to travel to a specific location. Humans that used a mixture of high and low individual autonomy achieved the highest performance levels (Walker et al., 2013). High individual robot autonomy allows humans to issue sequences of commands to the swarm as long as the human respects neglect benevolence (see Chapter II.2.2) and delays input of new commands in order to improve the swarm's performance (Walker et al., 2012a). Assuming that optimal human input times are calculated separately (Walker et al., 2012a), a recent algorithm determines the optimal command sequence to transition a swarm from an initial state to a desired goal state (Nagavalli et al., 2017). Future use of the algorithm is expected to assist humans in determining the order and timing of their control. One of the disadvantages of current human-swarm interaction studies is that humans must independently make all decisions about the swarm's future activity, often choosing from a finite set of available options.

Existing human-collective interaction research has demonstrated potential benefits of *shared-control*, in which the human and the collective each participate in the decision making process that affects the collective's current and future behaviors (Coppin and Legras, 2012; Crandall et al., 2017). Only two studies examine human interaction with collective decision making strategies (Coppin and Legras, 2012; Crandall et al., 2017). The first provides an interface to coordinate collaborative task allocation for human-collective defense of a simulated military installation (Coppin and Legras, 2012). The robotic collective and a group of humans shared a virtual pheromone map that ensured perceived environmental changes were rapidly shared with all team members. Global access to the virtual pheromone map enabled the humans and robots to make individual decisions about future tasks that implicitly considered all information available to the human-collective team. Human interaction was unnecessary for routine patrolling, but did improve the human-collective team's performance during the more difficult task of intercepting simulated intruders. An autonomy spectrum was introduced in order to describe multiple autonomy levels for each robot task and provide pathways connecting these levels to give humans the ability to quickly alter the system's autonomy during complex missions (Coppin and Legras, 2012). Task allocation is a combined decision making strategy (see Chapter II.1.1). Unlike best-of-*n* decisions, the collective decision results in a desirable distribution of robots executing various missions, rather than executing the same task.

The second study (Crandall et al., 2017) introduced the concept of shared control in human-collective teams making best-of-*n* decisions. The authors designed and evaluated a simulated hub-based colony using a honeybee nest site selection behavior (Nevai et al., 2010). Quorum detection and commitment mechanisms were used to transition the swarm from deliberation to execution following a decision, similar to Parker and Zhang's (2009) approach. It was suggested that a trade-off exists between the system's fault-tolerance and direct human control (e.g., increasing human control decreases system fault tolerance). A variety of control mechanisms and a preliminary user interface design were proposed. Unlike the robotic collective best-of-*n* studies in Chapters II.2.1 and II.2.2, Crandall et al. (2017) evaluated the collective in an environment containing obstacles, traps, and rough terrain. The collective made accurate decisions in many cases, but struggled to choose the best site available when a mediocre site was closer to the collective's decision making hub, which is consistent with previous observations of similar strategies (Valentini et al., 2017; Reina et al., 2015b). Methods were proposed to overcome the collective's limitations and inject global information into the human-collective decision making process, but the proposed mechanisms and interface have not been evaluated.

Function	Level of Autonomy		Action Selection Features
Acquire	Low:	Human directs exploration and evaluation of options.	Alter collective environment, individual exploration and evaluation behavior, and introduce or eliminate available options.
	High:	Collective identifies and evaluates options.	
Analyze	Low:	Human compares the identified options and reduces those considered.	Alter exploration activities. Alter direct (option-based) and indirect (environment-based) modulation.
	High:	Collective compares options, reducing the available set and converging to a single option.	
Decide	Low:	Human makes a decision for the collective.	(Explicit Only) Alter quorum threshold, trigger individual robot quorum detection, extend/reduce decision dissemination.
	High:	Collective makes an independent decision.	
Implement	Low:	Human initiates chosen activity.	(Explicit Only) Expedite, delay, and initiate collective transition to execution.
	High:	Collective initiates chosen activity.	

Table II.2: A summary of the different levels of collective autonomy and mechanisms that permit human control of a collective for each decision-making function described by Parasuraman et al. (2000).

II.3.1.2 Collective Action Selection Autonomy

The shared control approach (Crandall et al., 2017) is applicable to a wide variety of the previously described implicit and explicit action selection strategies. These strategies offer different levels of autonomy for the decision making functions: information acquisition, information analysis, decision and action selection, and action implementation (Parasuraman et al., 2000). Table II.2 summarizes the decision making functions, descriptions of collective autonomy levels, and features of collective action selection that permit different levels of human-collective interaction. Information acquisition is the gathering and filtering of sensory data, while information analysis integrates and interprets the acquired information and performs predictive analysis of various available actions, which results in a final set of options (Parasuraman et al., 2000). The decision and action selection function is the choice of one of the available actions identified during information analysis, while action implementation is the chosen action's initiation and completion (Parasuraman et al., 2000). Implicit and explicit collective action selection strategies complete each of these decision making functions in different ways. Features of each collective action selection behavior enable the human to adapt collective levels of autonomy, which has been shown to improve performance (Walker et al., 2013; Coppin and Legras, 2012).

The exploration and dissemination states in collective best-of-*n* strategies provide capabilities that are similar to information acquisition and analysis. Exploring robots interpret raw sensory data in order to identify and evaluate options, which provides rudimentary assessment of each option's value to the collective's actions (Valentini, 2017). The disseminating state in collective best-of-*n* strategies performs limited information analysis by comparing available options, promoting further evaluation of the most valuable options, and convergence to the best one. Although very few collective decision strategy studies consider more than two options (Valentini et al., 2017), the biological processes that inspire these strategies often reduce the number of available options to the most promising set (Seeley, 2010; Mallon et al., 2001). Implicit action selection accomplishes decision and action implementation concurrently with information acquisition and analysis. Exploring robots identify and execute available actions, while disseminating robots compare these actions in order to drive convergence to the best action. Human interaction with implicit action selection is; therefore, limited to these two functions even though all four decision making functions are accomplished.

Features of implicit collective action selection permit a variety of low to high levels of autonomy for information acquisition and analysis. At a low level of information acquisition autonomy, as indicated in Table II.2, the human directly controls the collective's available options, search space, option values, option costs, and search behaviors. The influence of these types of control decrease as more responsibility is shifted to the collective at higher levels of autonomy. Human influence during information analysis alters individual robot direct and indirect behavior modulation (see Chapter II.2.1) in order to alter the collective's comparison of available options. The identification and development of disparity resilience mechanisms in collective action selection strategies will enable humans to adjust actively the relative importance of discovery time, option cost, and option value in decision making strategies (see Chapter II.2.1). Further, manipulating the robots' rate of transition between exploration and dissemination enables the human-collective team to prioritize between information acquisition and analysis.

Explicit collective action selection separates information acquisition and analysis from the decision and action implementation functions. Individual robots initiate a collective decision when they detect that an available option has gained enough support, according to the collective's decision detection strategy (e.g., a change in the environment (Wessnitzer and Melhuish, 2003; Parker and Zhang, 2010) or quorum detection (Parker and Zhang, 2009)). Existing explicit collective action selection studies require the robots detecting that a decision has been reached to disseminate this information throughout the collective (e.g., using a decision message (Wessnitzer and Melhuish, 2003) or a separate decision state (Parker and Zhang, 2009)). Once the collective has reached a decision, it transitions to execution, which has been accomplished using timers (Wessnitzer and Melhuish, 2003; Parker and Zhang, 2009; Crandall et al., 2017), although sequential quorum detecting states are also possible (Parker and Zhang, 2010; Cody and Adams, 2017a).

The decision detection mechanisms characteristic of explicit action selection strategies control the duration of the first two decision making functions (see Chapter II.2.2). Introducing controls to manipulate decision detection mechanisms enables the human to initiate, delay, or expedite the collective information acquisition and analysis, which provides human control over the speed and accuracy of the decision making process (Parker and Zhang, 2009). Explicit action selection also permits manipulation of the collective's decision function. At low autonomy levels, the human decides between the options identified during information analysis without input from the collective. As the level of autonomy increases, the computer system narrows the number of options considered by the human (Sheridan and Verplank, 1978). Providing intermediate autonomy levels to the decision function in collectives will either be accomplished by reducing the options during information analysis or by altering the decision detection mechanism, possibly by lowering the quorum threshold, in order to permit agents to detect the desired level of support for a variety of options that the human considers in the final decision. The decision dissemination process is another possible mechanism that when used by the human provides an increase or decrease in the speed of transition between reaching a decision and executing that decision. Finally, different autonomy levels exist for the initiation of the chosen action. At low autonomy, the collective waits for a signal from the human to initiate the chosen action. At higher autonomy levels, the collective delays execution, which provides the human with time to intervene prior to execution or implements the action independent of human input.

Determining the appropriate level of autonomy for human interaction with collective action selection depends upon the collective best-of-*n* strategy, the types of options considered, and features of the collective's environment. Human operators will need to exert more control over the collective's behavior when the collective is likely to make an incorrect decision, or waste an excessive amount of resources or time when making a decision. Collective best-of-*n* decision making strategies are subject to the existing trade-off between decision speed and decision accuracy (Valentini et al., 2015) and the existence of undesirable biases (e.g., the distance from the decision making hub to a site (Reina et al., 2015b)). Finally, collective vigilance describes the collective's responsiveness to changes in the environment and human input (see Chapter II.2.1). A human who understands the strategy's expected decision speed, accuracy, disparity resilience, and vigilance will be better prepared to make decisions about appropriate levels of collective autonomy and intervene in order to improve the human-collective team's performance. Existing honeybee nest site selection inspired models (Crandall et al., 2017; Reina et al., 2015b) have limited resilience to distance and discovery time disparity, which means they are more likely to pick a closer mediocre site than a more distant optimal one (Crandall et al., 2017; Nevai et al., 2010). An informed human operator will reduce collective autonomy when a promising option is discovered late in the decision process or a promising option exists at the edge of the collective's search space. Similarly, understanding the collective's vigilance to environmental change or human input informs the human as to how much effort is required to alter the collective's decision process, when a change is deemed necessary. Greater collective vigilance implies that the collective responds rapidly to environmental changes and needs little human input to adjust to these changes. Lower vigilance means the human must increase the control exerted in order to overcome the collective's inertia to change. These metrics are only useful if the human receives enough information from the collective to understand the collective's state, its estimate of the environment, and its evaluation of available options. Providing information about the collective action selection process to the human is largely an open research area that is described further in Chapter II.3.2, but it critically influences human autonomy decisions and future performance of human-collective teams.

II.3.2 Collective Transparency

Human operators must understand the robot's current and future performance, reasoning, and intent in order to determine an appropriate level of trust (Chen et al., 2016). The ability of an interface to support a human's understanding of autonomous robots is known as *transparency* (Chen et al., 2017). Swarm transparency consists of assisting the human to understand the swarm's current and future physical states, which is challenging when communication limitations restrict the information the human receives from the swarm (Adams et al., 2018). The following questions address the challenges of collective transparency:

- What are the sources of uncertainty for collective action selection?
- How do decision making hubs improve collective action selection transparency?
- What visualization techniques improve collective action selection transparency?
- What are the communication requirements for collective action selection transparency?
- What techniques enable understanding of collective reasoning and past decisions?

Collective transparency extends swarm transparency to include intent, current and future collective decision making states, and explanations for past decisions (Adams et al., 2018). Additionally, collective explicit action selection introduces the need for human operators to understand sudden collective behavior changes. Collective decision making dynamics are complex and influenced by robot and environmental features that are not necessarily visible to the human (see Chapters II.2.1 and II.2.2). While collective information pooling behaviors can ease collective physical and decision making state estimation (Crandall et al., 2017), exploiting this behavior to best provide transparency remains an open research question.

The presented questions are examined by identifying known challenges in achieving swarm transparency and applying the presented topics to collective action selection transparency. Sources of uncertainty for human-collective interaction are introduced, which is used to frame directions for possible research in reducing uncertainty and improving transparency for collective action selection.

II.3.2.1 Sources of Uncertainty in Collective Action Selection

Both swarms and collectives are characterized by global behaviors that emerge from local interactions, which presents two immediate challenges to transparency. First, it is very unlikely that humans will receive physical state or command acknowledgment from every individual robot in either a swarm or a collective, so the emergent system state must be estimated from available information (Adams et al., 2018). Second, even with 100% complete information, humans have a limited ability to perceive emergent behavior in spatial swarms (Seiffert et al., 2015). Three different types of uncertainty exist in human-swarm interaction: physical, virtual, and compound (Hayes and Adams, 2014). Physical uncertainty includes the limitations humans face when interpreting noisy information about swarm behaviors. Human visual perception of biological swarm movements can distinguish swarm movements from random movements, even under noisy conditions, but this perception capability is less reliable than human perception of rigid body motion under similar conditions (Seiffert et al., 2015). Virtual state uncertainty represents the unobservable aspects of the swarm state, including leadership and information flow at the swarm and individual swarm member level. Finally, compound uncertainty is the simultaneous presence of both virtual and physical sources of uncertainty (Hayes and Adams, 2014).

Collective action selection requires extending the human-swarm interaction sources of uncertainty. Physical uncertainty in swarms is often reduced by exploiting the fact that robots are directly influenced by their neighbors (e.g., flocking (Goodrich et al., 2011; Walker et al., 2014) and incorporation of pattern formation (Kira and Potter, 2009)). Individual robots in collective action selection are less influenced by neighbor's proximity than by their internal states or opinions, which result in heterogeneous behaviors throughout the collective (Reina et al., 2015b; Scheidler et al., 2016). No human-collective studies have examined human perception of collective action selection, which means that predicting the required information to enable a human to understand the collective's physical state remains an open research question. It is expected that due to heterogeneous behaviors in collectives, humans will need more information from the collective's members and environment than is necessary for spatial swarms in order to understand the collective's physical state. Virtual uncertainty in collectives depends upon the human's ability to infer the collective's support for various options through observation. Perceiving the collective's decision making state will also depend upon how much of the environment and the collective's population is available to the human, either through direct observation or as viewed on an interface. Perception of the collective's state in a collective shortest path problem (Scheidler et al., 2016) requires the human to have access to the available paths and the relative number of robots traveling those paths over time. Acquiring this information through physical observation entails observing a large portion of the collective's space. Finally, compound uncertainty for collective action selection is complicated by the susceptibility of the decision processes to undesirable biases that alter individual and collective performance, either by changing collective opinion away from a desired choice, or by altering the time required for the collective to make a decision (see Chapter II.2.1). An example of the effect of environmental bias on collective action selection is a simple site selection scenario in which the human is unaware that an optimal site is partially occluded by an obstacle. Prior research showed that the collective is expected to either select the optimal site after a longer than expected convergence time or choose a suboptimal site (Reina et al., 2015b; Cody and Adams, 2017b). Finally, explicit action selection introduces collective behavior transitions, caused by internal robot state changes, that are not easily inferred from the physical states of the robots.

II.3.2.2 Reducing Uncertainty in Collective Action Selection

The two collective studies described in Chapter II.3.1 offer approaches to reducing uncertainty and providing transparency for human-collective interaction. The task allocation study (Coppin and Legras, 2012) established transparency with strong assumptions regarding the system's communication capabilities. Physical robot states and the virtual pheromone levels affecting individual decision making were readily visible to humans due to the universally shared virtual pheromone map (Coppin and Legras, 2012). Human perception of collective performance was error prone, even with near-perfect shared information between the humans and the collective (Coppin and Legras, 2012).

The site selection interface restricts human communication to a concentrated decision making hub (Crandall et al., 2017). This approach is advantageous if the robot interactions are well-mixed (see Chapter II.2.1) and restricted to the hub, as sampling the states and opinions of robots within the hub will enable an estimate of the collective's overall exploration and decision making behaviors (Crandall et al., 2017). Further, if the hub is smaller than the collective search space, the communication requirements throughout the collective's search space are reduced. A radial display, focused on the decision making hub, was proposed to provide information about the relative distribution of the robots within the environment as well as values of known options (Crandall et al., 2017). This information provides the human with estimates of the collective's current state and indicates projected states for individual robots. The use of the decision making hub to enhance collective transparency introduces a well known trade-off between the number of observable robots and the length of observation time required to make an accurate state estimation (Brown et al., 2016). State estimation accuracy is limited by the size of the decision making hub and the expected number of robots within the hub at any given time. Increasing the size of the hub and the available robots decreases the time to determine a state estimation, increases the estimation's accuracy, and increases the communication requirements necessary to receive information from the hub. While the interface is promising, it does not readily apply to action selection strategies with mobile collectives (e.g., collective target tracking (Wessnitzer and Melhuish, 2003)) or those whose decision making hubs are similar in size to the collective's search space (Parker and Zhang, 2009, 2010).

Each of these strategies requires representations of the collective location within the environment, along with a similar abstraction of the relative support for each of the collective's options. A recent study evaluated different visualization techniques with flocking, dispersing, and aggregating swarms and determined that humans were equally capable of understanding current and future swarm state using either a full representation of the swarm (e.g., consisting of all individual member positions) or a combination of an enclosing ellipse with a center point (Walker et al., 2016). Collective action selection transparency is likely to benefit from a similar enclosing ellipse, along with an overlaid visual representation of the collective's decision making state (e.g., a collective opinion centroid that moves between available opinions).

The previously described approaches to reducing uncertainty in collective action selection primarily address the system's current physical and decision making states, but offer little information about the rationale for a given state, projected future states, and reasoning behind prior decisions. Implicit and explicit action selection offer different possible solutions. An individual robot in collective implicit action selection does not necessarily know that it is executing the best available action or have any other information regarding the collective's decision making state. A simple classifier has been shown to be effective for determining spatial swarm states from limited observation of a small number of individual robots (Brown et al., 2016). A similar approach, in which a collective's population is sampled based on their location in the environment, is likely to improve the collective's physical and decision making state estimates when robot behaviors between options are observably different. Although explicit action selection is more complicated than implicit action selection, due to its collective behavior transitions (Adams et al., 2018), its use of decision detection mechanisms provides individual robots with basic estimates of collective state (Parker and Zhang, 2009). Under the assumption of well-mixed robot interactions, a sampling of quorum information from robots will provide an estimate of option strength for the current decision. Once a decision is detected, information regarding the decision is communicated throughout the collective in order to prepare it to transition behaviors, which provides individual robots with information regarding the previous decision outcome. This information is valuable, in that it offers the ability to estimate option comparison and decision trends, but it does not capture information related to a collective's failure to select a particular option. Running estimates of cost and discovery time disparity between options will likely provide some insight into these situations, but this capability is not described by current collective action selection strategies. The degree to which individual robot behaviors and capabilities, as well as communication requirements must be altered in order to achieve reasonable running estimates of collective disparity resilience remains an open research question.

II.3.3 Collective Control

Control mechanisms are the focus of several human-swarm interaction evaluations, as indicated by recent surveys (Coppin and Legras, 2012; Kolling et al., 2016), but have only recently been applied to human interaction with collective action selection strategies (Crandall et al., 2017). Interactions with physical robot states, such as reducing or increasing robot exploration behaviors, are likely to be similar to mechanisms used to control spatial swarm behaviors (e.g., placing attracting or repelling beacons (Kolling et al., 2013; Crandall et al., 2017) or adjust exploration parameters (Kira and Potter, 2009; Crandall et al., 2017)). The collective decision strategies rely on positive feedback mechanisms that are related to the strength of available opinions and the existing bias within the environment (Valentini, 2017). The strength of the positive feedback; therefore, differs depending upon which option the human intends to promote or reduce within the collective. Several open research questions remain for human control over collective action selection:

- What control mechanisms enable human control over the collective action selection's physical and decision making states?
- How do communication considerations alter control mechanisms for collective action selection?
- How does the concept of neglect benevolence apply to collective action selection?

A recent human-swarm interaction survey provides a description of the current control mechanisms used to enable humans to influence spatial swarm behaviors (Kolling et al., 2016). Some of these mechanisms have been recently considered, though not evaluated, for human-interaction with collective action selection in a site selection scenario (Crandall et al., 2017). The existing human-swarm and human-collective control mechanisms are applied to the collective action selection strategies described in Chapters II.2.1 and II.2.2. Consideration is given to the practical use of the proposed mechanisms, which depends upon human understanding of the collective's underlying dynamics and the communication requirements for implementing the desired controls. Finally, the concept of neglect benevolence is applied to collective action selection in order to identify risks associated with poorly timed human input. Table II.3 provides a summary of the existing spatial swarm control mechanisms that are considered for application to future interaction with collective action selection strategies.

Control	Swarm Description	Collective Action Selection
Parameter Setting	Adjust parameter values to change swarm behavior.	Adjust exploration and dissemination *, quorum sensing, and disparity resilience parameters.
Behavior/Algorithm Selection	Direct execution of preprogrammed behaviors.	Direct exploration, alter option support, alter decision detection.
Environment Influence	Adjust features in the environment to cause behavior change.	Alter collective search space and modify option evaluation *.
Playbook	Develop and validate a plan composed of human-robot understood plays (not yet implemented with swarms).	Switch collective behaviors *; enable and disable considered complex behaviors.
Leader Selection	Exerts influence through proxy swarm members, or unique external agents.	Guide exploration; manipulate deliberation; human-operator as collective member.

Table II.3: A combination of control strategies for swarms (Kolling et al., 2016) and collectives (Crandall et al., 2017; Coppin and Legras, 2012) are extended to the collective implicit and explicit action selection strategies. Each control strategy considered in the current site selection study is identified by the * symbol.

II.3.3.1 Parameter Setting

Both spatial swarm and collective action selection behaviors typically emerge as a result of complex interactions between individual robots (Brown et al., 2016). Humans with the ability to manipulate various parameters can influence group behaviors by altering how the robots interact with each other and the environment (Kolling et al., 2016). This control mechanism is most effective when the human has global access to the spatial swarm (Kira and Potter, 2009), but it is typically of limited use when the human lacks global access to the swarm (Kolling et al., 2016) or lacks understanding of the underlying system dynamics (Crandall et al., 2017). The use of parameter settings in collective action selection scenarios requires that the human understand the collective's physical and decision making behavior dynamics. Parameter setting was proposed for site selection scenarios in order to control the robot's direction of exploration when leaving the decision making hub and the transition rate between exploration and dissemination states (Crandall et al., 2017). The advantage of the decision making hub is that every robot periodically returns to the hub during the decision making process associated with most collective action selection strategies (see Chapter II.2.1). Parameter adjustments that are persistently transmitted to the decision making hub are likely, but not guaranteed, to affect the parameters of most robots within the collective. Several other parameters for consideration include random walk parameters for exploration (e.g., range and duration), disseminating state modulation parameters (e.g., option evaluation criteria), and quorum threshold values, which determine the duration and accuracy of the decision process in explicit action selection strategies. Additionally, if the human can manipulate a collective's disparity resilience (see Chapter II.2.1), it is possible to dynamically adjust the collective's decision making priorities between option value, option cost, and option discovery. Adjusting any of the described parameters changes how the collective explores its environment and compares the discovered options, but there are two primary disadvantages to this approach. First, humans are required to understand the underlying system dynamics in order to effectively alter parameters (Crandall et al., 2017). Second, if parameter adjustments are not applied to the entire collective, then heterogeneous behaviors are introduced. These disadvantages are a primary reason that parameter adjustment is often limited to system design (Kolling et al., 2016).

II.3.3.2 Behavior and Algorithm Selection

Behavior and algorithm selection is a control mechanism that requires humans to select a subset of members within the swarm, or collective, and direct them to either change their behaviors or alter the algorithm used for interaction with robots and the environment (Coppin and Legras, 2012; Kolling et al., 2016). This control mechanism is common in swarm interactions, although it is typically reserved for scenarios in which the human desires to change the behavior of a subset of the swarm (Kolling et al., 2013; Walker et al., 2013). Behavior selection can potentially be very useful in collective action selection (e.g., different sites or paths). The information pooling that occurs during collective action selection means that if agents are commanded to support a known option, investigate a location for possible options, or abandon an option, the influence on the collective decision strategy is likely to expand beyond the selected individual agents. One example of this type of interaction is the introduction of a new path in a collective shortest path problem. A new path discovered late in a collective action selection scenario is unlikely to be chosen by the collective (Scheidler et al., 2016). The human can encourage the collective to select the new path by commanding robots, possibly within the decision making hub, to abandon their current path and investigate the new path. A combination of these activities reduces the strength of the collective's support for the former path and exploits the system's positive feedback to promote the new path. This mechanism still requires the human to understand the collective's decision making dynamics, but it does not introduce inconsistent parameters and is unlikely to require the human to interact with most of the collective's individuals.

II.3.3.3 Environmental Control

Environmental control mechanisms refer to any physical or virtual manipulation of environmental features that cause a swarm, or collective, to react in a predictable manner (Kolling et al., 2016). Examples of environmental control include the use of virtual pheromones to alter individual behavior decisions (Coppin and Legras, 2012) and the placement of beacons to cause behavior changes to robots within a beacon's range (Kolling et al., 2013). The use of attracting and repelling environmental controls was recently recommended as a means of encouraging and discouraging exploration in certain areas of the collective's environment (Crandall et al., 2017). These strategies are often persistent (Kolling et al., 2013) or semi-persistent (Crandall et al., 2017; Coppin and Legras, 2012), which potentially reduces the number of human interactions required to result in the desired behavior. Physical changes to the collective's environment requires either a human or another robot to travel to a specific location and alter it in a desired manner. This type of interaction can be useful when humans are in close proximity to the collective (Alboul et al., 2008), but it is often undesirable to require physical adjustment to the environment when remotely controlling a collective (Kolling et al., 2016). Virtual environment manipulation introduces the challenge of distributing and sharing virtual environment information and is often achieved using global access to an environmental map (Kira and Potter, 2009; Coppin and Legras, 2012). Available communication latency, bandwidth and range are likely to restrict the human-collective team's ability to globally share this type of information. These challenges increase the likelihood that individual robots will operate with outdated environmental information. While it is conceivable that the human-collective interface can provide virtual environment updates to the robots entering or exiting the decision-making hub, the limitations of this approach's effectiveness and communication requirements have not been evaluated.

II.3.3.4 Playbook

Playbook control consists of a set of *plays* shared by the human and robotic team members that represent plans for achieving goals during deployment (Miller et al., 2005; Miller, 2012). Playbook's success in tournament style games, such as RoboFlag (Campbell et al., 2003; Parasuraman et al., 2003; Chamberlain et al., 2003) has inspired its development for swarms (Coppin and Legras, 2012; Crandall et al., 2017), although it has not yet been implemented for collectives. This control mechanism was developed in order to enable a human to build, validate, and use a set of plans to alter the behavior of a multi-robot team. Playbook has also been considered to facilitate collective switching between implicit action selection behaviors (e.g., site selection) and other typical spatial swarm behaviors (Crandall et al., 2017). Alternatively, Playbook enables a human to describe complex actions that are provided to the collective as an available set of plays for consideration by the human-collective team. This approach can potentially enable very complex collective actions, but also aggravates the existing challenges of human interaction with explicit collective action selection by introducing numerous collective behavior transitions.

II.3.3.5 Leader Selection

Leader selection refers to the human's interaction with one, or more, robots to exert influence over the swarm or collective (Kolling et al., 2016). Leader selection mechanisms have been developed that use unique external robots (e.g., Goodrich et al., 2011) or selected members of the swarm (e.g., Penders, 2007). These control mechanisms rely on direct interaction between leaders and individual robots in order to influence nearby robots to respond in predictable ways (Walker et al., 2014; Goodrich et al., 2011; Jung and Goodrich, 2013; Jung et al., 2013). Two human-collective interaction opportunities exist for this type of control mechanism. First, during exploration, the use of human controlled leaders can be beneficial in directing and guiding exploration behaviors. Second, leader selection strategies can enable human controlled robots to exert greater influence over decision making. This control mechanism requires human control throughout the collective area (Kolling et al., 2016). Despite this disadvantage, the control mechanism can be beneficial in heterogeneous collectives and is similar in concept to the unique agents proposed in spatial swarm scenarios (Goodrich et al., 2011). Special agents with different interaction parameters were used to influence spatial swarm behaviors according to the human's intent. These unique agents either repelled or attracted nearby robots, which enabled the human to guide the swarm's behavior. Applying this concept to the dissemination state enables humans to exert greater influence over the decision strategy, without necessarily overriding the collective's decision making process. An excellent example of the usefulness of this approach is a large human-collective team in which humans exert control at two levels. At the collective level, the human influences the collective's overall exploration and decision making processes. At the individual level, humans control either members of the collective or more capable individual robots, and participate in the collective decision making process. Human controlled robots can be given greater priority in the decision process, although this approach introduces heterogeneous decision influence that must be evaluated.

II.3.3.6 Neglect Benevolence

Human-swarm teams often perform better when humans respect the swarm's neglect benevolence and delay input in order to permit the swarm to stabilize in response to a prior command (Walker et al., 2012a). Prior research has demonstrated the importance of neglect benevolence in human-interaction with spatial swarms (Walker et al., 2012b,a) and methods have been proposed for calculating optimal human input times to account for this phenomenon (Walker et al., 2012a). This concept has yet to be applied to collective action selection. Typically, neglect benevolence implies that the human input disrupts the swarm's current behavior (e.g., when a flocking swarm receives a new heading (Walker et al., 2012b)) and the swarm requires time to stabilize prior to receiving the next command (Kolling et al., 2016). Stable states in collective best-of-*n* strategies are achieved when a majority opinion (e.g., Reina et al., 2015b) or consensus (e.g., Valentini et al., 2015) exists that supports one of the options available to the collective. Delaying human input when the collective is deliberating between known options is unlikely to be beneficial. Collective support for individual available options increases over time as a result of the decision strategy's direct and indirect modulation mechanisms (see Chapter II.2.1). A human who intends to alter the collective's chosen option will only have a larger population of robots supporting an undesirable option if influence is delayed during the collective's deliberation of options. Despite the apparent absence of neglect benevolence during collective deliberation, this dissertation proposes two opportunities for examining neglect benevolence in collective action selection. First, strategies that find options within an environment during decision making (e.g., Parker and Zhang, 2009; Reina et al., 2015b) are likely to perform better if the human does not increase or decrease collective support for options too early in the discovery process. A human who increases support prematurely for an option reduces the number of robots available for exploration and increases the challenge of altering the collective's decision, if a better option is discovered later. Decreasing option support too soon introduces the problem of rebuilding support if it is found to be the best available option after all. The second application of neglect benevolence is specific to collective behavior transitions in explicit action selection strategies. Once a collective has reached a decision and begins transitioning to a new behavior, the collective is unstable and the robot population is split between deliberation and action implementation states. Human input during this period is likely to cause the collective to fragment in a manner similar to that described for spatial swarms (Walker et al., 2012a), as deliberating robots are left behind by those that have been induced to implement the action too soon.

II.3.4 Complex Organizations

Collective action selection strategies are readily applicable to large heterogeneous and composite organizations for two primary reasons. First, they require a population of dedicated individual robots to explore the collective search space and conduct lengthy deliberations between available options (see Chapter II.2.1), which often deviates from the collective's primary task, as in the biological decision processes (see Chapter II.1.2). Future systems will likely benefit from a similar heterogeneous population, although research into collective and swarm heterogeneity is limited (Dorigo et al., 2013; O'Grady et al., 2009; Ducatelle et al., 2011; Kengyel et al., 2015). Second, implicit and explicit action selection are capable of high levels of autonomy that will enable humans to monitor and influence larger organizations comprised of multiple collectives, which has been proposed and evaluated in highly autonomous spatial swarm behaviors (Brown et al., 2016), but has yet to be evaluated in human-collective interaction. The following questions are considered:

- What is the effect of heterogeneity on collective action selection strategies?
- How can collective action selection be incorporated as the interactive decision making component of a

larger heterogeneous organization?

- What challenges exist for human control of multiple collectives conducting action selection within a shared space?

Prior work with heterogeneous robotic swarms and human interaction with multiple swarms is considered and applied to challenges for future human-collective teams. The effect of heterogeneity on collective best-of-*n* decisions and human-collective interaction is considered. Finally, human control of multiple collectives conducting action selection is discussed.

II.3.4.1 Heterogeneous Collective Action Selection

Most robotic swarm studies assume a homogeneous swarm population, which ensures the swarm's behavior is robust to individual robot failure, but ignores the potential performance benefit of more diverse populations (Kengyel et al., 2015; Dorigo et al., 2013). A few studies have examined physical and behavioral heterogeneity in swarm robotics (Dorigo et al., 2013). The former refers to collectives composed of robots with different sensing, mobility, and interaction capabilities (e.g., O'Grady et al., 2009; Ducatelle et al., 2011), while the latter involves physically homogeneous robots that exhibit different behaviors (e.g., Kengyel et al., 2015). Evaluation of heterogeneity in collective action selection is important, since it can enable incorporating different robotic capabilities, and because human control mechanisms, such as leader selection (see Chapter II.3.3), can induce heterogeneous behaviors.

Behavioral heterogeneity within the population of robots conducting consensus achievement has been evaluated in flocking (Ferrante et al., 2012), aggregation (Kengyel et al., 2015), and collective best-of-*n* decision making (Reina et al., 2015b) using homogeneous robots that differ in their execution of one or more behaviors. The aggregation study examined different robot mobility patterns to improve the collective's ability to achieve consensus supporting the best aggregation site (Kengyel et al., 2015). Another study examined heterogeneous flocking, where all robots executed two of three flocking behaviors, but only a subset of the swarm's population were capable of aligning with their neighbors (Stranieri et al., 2011). Finally, behavioral heterogeneity was examined in a site selection scenario in which individual robots used heterogeneous, randomly distributed site acceptance thresholds to determine whether a site was acceptable or not (Reina et al., 2015b). The heterogeneous model was compared to a homogeneous model in which all robots shared the same site acceptance threshold. Both models achieved similar accuracy in a number of scenarios, but the homogeneous model outperformed the heterogeneous model when choosing between an optimal site with a higher cost than a mediocre site.

Physical heterogeneity is expected to be important to future human-collective action selection. The introduction of robots with differing sensing, mobility, and interaction capabilities into the decision making process was briefly considered in Chapter II.3.3.5 by proposing the addition of more capable, human-controlled individual robots. Additionally, complex collective search spaces, such as those with portions of the environment that are either difficult to navigate or highly structured, are likely to require robots with differing mobility and sensing capabilities. Either scenario introduces similar problems to those demonstrated by the heterogeneous site selection study (Reina et al., 2015b). Naturally, faster robots with greater interaction capabilities will bias the system, perhaps intentionally, towards their supported option. Overcoming cost disparity becomes easier if the more capable robots identify higher cost optimal options, but the opposite is true if the same option is discovered only by less capable agents. Improving a human-collective's disparity resilience in these situations requires that control mechanisms and transparency techniques account for heterogeneous robot influence.

Explicit action selection in biological colonies is often allocated to a group of scouting insects that explore for options and deliberate between them, while the majority of the collective continues colony activities (Seeley, 2010; Mallon et al., 2001). Researchers have reproduced behaviors similar to *Apis mellifera* honeybees and *Temnothorax albipennis* ants (see Chapters II.1.2.3 and II.1.2.4), but have yet to employ these strategies as a means of coordinating the activities of a larger heterogeneous collective. The interaction of physically and behaviorally distinct swarms of robots that cooperate to accomplish complex tasks has been examined in one series of studies (Dorigo et al., 2013; O'Grady et al., 2009; Pinciroli et al., 2010; Mathews et al., 2012; Pinciroli et al., 2009; Ducatelle et al., 2011). The Swarmanoid concept (Dorigo et al., 2013) consists of distinct swarms of ground, air, and manipulator robots that are evaluated for complex tasks. Most Swarmanoid studies use aerial robots, called *eye-bots*, to inform and guide ground robots, called *foot-bots*, in completing a collective task, such as aggregation (O'Grady et al., 2009), foraging (Pinciroli et al., 2009), and coordinated movement (Pinciroli et al., 2010). While the eye-bots guided the foot-bots' behaviors, they also responded to changes in the foot-bot behaviors, which resulted in mutual adaptation between these distinct groups (Pinciroli et al., 2010). One study used eye-bots to prioritize foraging behaviors of foot-bots in different areas through the use of parameter setting techniques, similar to those described for human-swarm control in Chapter II.3.3.1 (O'Grady et al., 2009). Foot-bots aggregated to locations defined by stationary eye-bot locations. Meanwhile, each eye-bot controlled its individual quota of foot-bots at its location by altering the aggregation parameters of the foot-bots within its communication range (O'Grady et al., 2009). The extension of this concept to collective action selection is straightforward. A group of distinct, decision making systems either constantly monitors the behavior of the rest of the collective, or emerges as a result of an environmental cue and begins the collective action selection process that will ultimately alter the entire collective's behavior in a manner similar to honeybee nest site emigrations (Seeley, 2010).

Increasing the complexity of collective behaviors by introducing distinct behavioral or physical subgroups of robots presents several implications to human-collective interaction. First, human control mechanisms already induce heterogeneity into the system. Leader selection strategies (see Chapter II.3.3.5) introduce distinct agents that often have unique movement and interaction capabilities (e.g., Goodrich et al., 2011; Jung and Goodrich, 2013). Parameter setting control mechanisms that alter a portion of the collective (e.g., robots in a decision making hub (Crandall et al., 2017) or near an environmental beacon (Kolling et al., 2013)) introduce heterogeneous parameters that result in distinct behavior groups within the collective. Second, human interaction with heterogeneous collectives requires that different collective behaviors be transparent to the human and receptive to the human's commands. Third, the calculation of optimal human input time (Walker et al., 2012a) necessary to prevent undesirable, unstable collective states (e.g., fracturing) must account for interactions between the decision making group and the remainder of the collective. Fourth, the added complexity of heterogeneity is expected to increase the human's workload (Kolling et al., 2016), who must exert additional effort to understand and interact with disparate behaviors.

II.3.4.2 Multiple Collectives

Highly autonomous swarm and collective behaviors provide an opportunity for large groups of decentralized robots to be controlled by a single human, if the underlying dynamics of the groups' behaviors possess stable attractors (Brown et al., 2016). Human-swarm interaction researchers have frequently demonstrated that humans are capable of controlling a single swarm (e.g., Walker et al., 2012b; Goodrich et al., 2011; Goodrich et al., 2012). Multiple swarms that use highly autonomous behaviors, such as flocking (Couzin et al., 2002) can be controlled by a single human with intermittent observation and influence over each swarm (Brown et al., 2016). The high levels of autonomy in collective action selection strategies (see Chapter II.3.1) are expected to enable similar human control over numerous collectives.

Two immediate advantages exist for human supervision of multiple collective action selection processes. First, multiple collectives in the same search space enables concurrent best-of-*n* decision processes to identify multiple valuable options within the environment. Second, individual collectives can be designed with a specific maximum range that reduces the influence of cost and discovery disparity on the collective's decisions (see Chapter II.2.1). Increasing the number of collectives within a human-collective team also introduces numerous problems that have yet to be explored. The possibility of collective fragmentation, as a result of either false quorum detection (see Chapter II.2.2) or poorly timed human input (see Chapter II.3.3.6) becomes more problematic, especially if the human must recover from the simultaneous fragmentation of multiple collectives. Redundant selection of a single site, which results in two collectives moving to the same site and greatly increasing congestion, is another problem that must be examined.

II.3.5 Summary

This chapter examined collective best-of-*n* decision making strategies in biological and robotic collectives and applied these behaviors to the existing state of the art in human-swarm interaction. Prior work has indicated that human-collective shared control informs human decisions and promotes the development of fault tolerant

Topic	**Research Question**
Autonomy	What levels of autonomy are available for human interaction with collective action selection? What collective action selection features enable low, intermediate, and high levels of autonomy for adaptive human-collective interaction? What evaluation criteria assist designers in determining an appropriate level of autonomy for different collective action selection strategies?
Transparency	What are the sources of uncertainty for collective action selection? How does the presence of decision making hubs improve collective action selection transparency? What visualization techniques enable improved collective action selection transparency? What are the communication requirements for collective action selection transparency? What techniques enable human operators to understand collective reasoning and past decisions?
Control	What control mechanisms enable human control over the collective action selection's physical and decision making states? How do communication considerations alter available control mechanisms for collective action selection? How does neglect benevolence apply to collective action selection?
Complex Organization	What is the effect of heterogeneity on collective action selection strategies? How can collective action selection be incorporated as the interactive decision making component of a larger heterogeneous organization? What challenges exist for human control of multiple collectives conducting action selection within a shared space?

Table II.4: Open research questions for human interaction with collective action selection strategies.

team operation (Crandall et al., 2017). This chapter extends the discussion in order to identify open research questions associated with collective action selection strategies for human-collective teams. These questions, summarized in Table II.4, address challenges inherent to the autonomy, transparency, and control of single, multiple, and heterogeneous collectives in future human-collective teams.

Collective action selection strategies are capable of operating at several different levels of autonomy (see Chapter II.3.1), with many possible mechanisms of injecting human influence (see Chapter II.3.3), but the human's choice of the appropriate level of autonomy and amount of input into the collective's decision making process depends upon the human's understanding of the collective's behavior (see Chapter II.3.2). The field of human-swarm interaction has focused more on exerting human control over the swarm, rather than on improving swarm transparency (Kolling et al., 2016), although recent results have narrowed this gap (e.g., Adams et al., 2018; Walker et al., 2016). Arguably, transparency in human-collective interaction is more important than in human-swarm interaction, since collective action selection strategies are most beneficial when the human is able to delegate some of the human-collective decision making functions to the collective. Humans will need to rely on the collective's decision making capability when the human is either unable to attend to every aspect of the decision making process, or has insufficient information about the collective's state and environment (Crandall et al., 2017). Human-collective performance will suffer if the human is unable to trust the collective to accomplish some, or all, of the team's decision making functions (Chen et al., 2016). During human-collective interaction, the human must be able to quickly determine the presence, cause, and resolution of a deviation between the collective's expected and desired future states. Future human-collective interaction research must focus on improving the human's understanding of the collective's decision making process and limitations. This capability enables the human to make informed choices regarding appropriate levels of autonomy, the type of control to exert, and the amount of influence to introduce.

Determining the collective's expected future state requires the human to understand the collective's current state and expected performance. Challenges associated with estimating, communicating, and portraying the collective's current physical and decision making states were summarized in Chapter II.3.2.2. Establishing collective action selection performance expectations and portraying these to the human remain important directions of future research. As discussed in Chapter II.2.1, collective best-of-*n* decision strategies differ in their expected accuracy and decision time (Valentini et al., 2014, 2015), as well as in their resilience to undesirable biases (Reina et al., 2015b, 2016; Cody and Adams, 2017b). Currently, there are very few comparison studies that offer insight into performance differences between collective best-of-*n* decision strategies (Valentini et al., 2015) or into collective disparity resilience (Reina et al., 2015b; Cody and Adams, 2017b). The human will require estimates of expected performance, based on these metrics, as well as the current biases within the collective's environment. A human who understands the collective's expected performance and the biases currently influencing the collective will be able to estimate expected future collective states and the cause of deviation from a desired state. This understanding can inform the human's decision to adjust the collective's level of autonomy or exert control over the collective process. Collective action selection strategies that use decision making hubs are likely to ease the challenges of transparency by enabling the interface to sample information from robots within the hub (Crandall et al., 2017), depending upon communication constraints (e.g., bandwidth, range, and latency). Understandably, increasing the information requirements from the collective introduces new challenges that must be examined in future studies. The quality of interface estimates depends upon the mixing of the robot population within the hub, the size of the hub, the robots' sensing ability, and the available communication structure.

Resolving a difference between future desired and expected collective states requires the human to understand the likely effect of the timing, type, and strength that available control mechanisms will have on the collective's decision making process. Providing the human with feedback related to the expected effect of human input is a problem that is common to human interaction with swarms (Kolling et al., 2016) and collectives. Complicating this feedback is the direct and indirect behavior modulation that determines robot state transition rates between states supporting various options. These modulations, which result from biases within the environment and the dissemination activities of subgroups of the collective population (see Chapter II.2.1) increase the difficulty of determining how much human control is necessary to alter support for an available option or neutralize an undesirable bias affecting the system. Predictive displays that provide the expected affect of human control inputs, similar to those used in human-swarm interaction in coordinated movement tasks (Walker et al., 2012b; Nagavalli et al., 2017), will provide the human with effective options.

Providing collective transparency to the human in scenarios with multiple or heterogeneous collectives becomes both difficult and critical, since the human is expected to have less ability to focus on any one collective for an extended time period. The transparency requirements for making autonomy and control decisions were briefly discussed, but future work must establish the most efficient means of collecting and visualizing information about each collective's expected performance and possible response options in order to alter each collective behavior in a desired way.

Chapter III

Robotic Collective Action Selection

A novel approach is presented for robotic collective action selection, in which a collective chooses the best option within a search space despite existing environmental bias, implements an action associated with the choice, and returns to a decision making state. The approach integrates a value-sensitive decision making process and permits the collective to operate from a centralized decision making hub. Value-sensitivity enables the collective to delay making a decision between unacceptable options (Pais et al., 2013). The decision making hub serves three important purposes. First, restricting robot interactions to a hub improves opinion mixing within the population. Second, the collective operates without being confined to an enclosing arena (e.g., Parker and Zhang, 2009; Valentini et al., 2014; Mermoud et al., 2014), and finally the hub provides human operators access to the collective, which is explored in Chapter IV.

This chapter focuses on the proposed explicit collective action selection process, which seeks to address the first four research questions presented in Chapter I.

R_1: How can a collective make efficient and accurate decisions when the problem space is subject to negative environmental bias?

R_2: How can *well-mixing* dynamics be adjusted to improve collective decision accuracy and efficiency?

R_3: How responsive are collective best-of-*n* decision making processes to changes in option values?

R_4: How can a collective transition between decision making and task implementation behaviors?

Two existing models, which are inspired by *Apis mellifera* honeybee (Seeley et al., 2012) and *Temnothorax albipennis* ant (Mallon et al., 2001) emigration, are combined and extended (Parker and Zhang, 2010; Reina et al., 2015b). The first model is referred to as the *original model* and represents Reina *et al.*'s (2015b) collective best-of-*n* decision making model. The original model accepts option discovery during the decision making process, makes value-sensitive decisions, and relies on the use of a decision making hub. The model faces its hardest decisions when it must choose a higher cost, optimally valued option over a mediocre option with lower cost. These problems are important, but under evaluated in collective decision making strategies (Valentini et al., 2017). Additionally, the original model does not detect when a decision has been reached or alter the collective's behavior. The second primary influence is Parker and Zhang's (2009; 2010; 2011) quorum detection-based model that enables collective behavior transitions, but does not use a value-sensitive decision making process or a decision making hub.

The chapter begins with a formal definition of the collective action selection problem, followed by an introduction to the collective and individual behavior models that address the four research questions. Finally, the results of four experiments comparing the new model performance to the original model are presented and discussed.

III.1 Formal Problem Definition

The formal definition merges features of the problems addressed by Reina *et al.* (2015b) and Parker and Zhang (2009). The agent descriptions are similar to those used by Kouvaros and Lomuscio (2016), which described agent states for a majority voting collective best-of-n decision making model (Valentini et al., 2015).

A single collective action selection problem requires a collective to find a set of options located within the environment, identify the best single option according to a given value function, and execute an action associated with the chosen option. This problem restricts the number of actions to the number of available options within the environment. Action execution is either implicit or explicit. Implicit collective action selection problems are those in which the collective must choose between the actions used to assess the available options (e.g., converge to the shortest path). Explicit collective action selection problems are those in which the collective chooses an action associated with an option that is not in the set of its assessment actions.

Definition 1. Collective Action Selection Problem: Each collective action selection problem is a tuple $CAS = \langle \Phi, A, Act \rangle$

- $\Phi = \{\phi_1, .., \phi_n\}$ is a set of n options located within an environment.
- $A = \{a_1, .., a_N\}$ is a collective comprised of N agents.
- $Act = Act_{assess} \bigcup Act_{choose}$ is the set of required agent actions, where Act_{assess} is the set of option assessment actions and Act_{choose} is the set of actions the collective is choosing between. When $Act_{choose} \nsubseteq Act_{assess}$, the problem is an explicit action selection problem. When $Act_{choose} \subseteq Act_{assess}$, the problem is an implicit action selection problem.

The set of options, Φ, contains the observable features within the environment (e.g., paths, sites, resources, or environmental conditions) associated with Act_{choose}. Options are assumed to be distinguishable from each other and from the remaining environmental features.

Definition 2. Option: Each individual option is a tuple $\phi_i = \langle f_i, v_i, cost_i, loc_i \rangle \forall i \in \{1, .., n\}$, where

- $f_i \in \{true, false\}$ is $true$, if the option has been identified by an agent and $false$ otherwise.
- $v_i \in [0,1]$ is the value of ϕ_i.

- $cost_i \in [0,1]$ is the cost required to find and assess ϕ_i (e.g., time and resources).
- $loc_i \in \mathbb{R}^{dim}$ uniquely identifies ϕ_i (e.g., location), where dim is the necessary dimension.
- ϕ_0 denotes the absence of an option, where $\phi_0 = \langle f_0 = false,\ v_0 = 0,\ cost_0 = 0,\ loc_i = \emptyset \rangle$.

The definition of an option is intended to be broadly applicable to a variety of different alternatives (e.g., paths or resources), the experiments in this dissertation focus on the comparison between sites, which are specific locations in the environment identified by a location in either two or three dimensional space.

Definition 3. Site: A site is a type of option that is assumed to be uniquely identified by its location, $loc_i \in \mathbb{R}^{dim}$, where dim represents the dimensions, two or three, and is dependent on the intended search space.

Individual agents within the collective are assumed to be capable of finding and assessing options, maintaining an opinion associated with an option $\phi_i \in \Phi^{found}$, interacting with neighboring agents, and executing an action associated with an option. Agents alter internal state as a result of interactions with neighboring agents, or as a result of assessing or abandoning options in the environment.

Definition 4. Agent: Each agent, for $m \in \{1,..,N\}$, is a tuple, $a_m = \langle h, s, checkQ, t^I, t^L \rangle$ where

- $h : \phi_i \rightarrow v_{mi} \in [0,1]$ maps an option to its estimated value for a_m; $h(\phi_0) = 0$.
- $Q_m = \{q_1, ..q_k\}$ is a queue of neighbor opinions of length $k \in \mathbb{N}$, where $q_j \in \{true, false\} \forall j \in \{1..k\}$ A $true$ value represents an interaction with an agreeing agent.
- $s \in S = \{\langle \phi_i, v_{mi}, Q_m, l, act_m \rangle : v_{mi} \in [0,1], l \in \{true, false\}, act_m \in Act\}$, defines a local state where $l \in \{true, false\}$ represents that the agent is latent ($true$) or interactive ($false$).
- $checkQ : Q_m \rightarrow \{true, false\}$ determines the presence of a quorum for a_m's current state.
- $t^L : S \times \Phi$ is a latent transition function, where $t^L(s, \phi_j) = s' \ \forall j \in \{0,..,n\}$.
- $t^I : S \times S^{neighbor} \rightarrow S$ is an interactive transition function, where $S^{neighbor}$ contains the states of a_m's neighboring agents and $t^I(s, S^{neighbor}) = s'$.

A collective solves an action selection problem by choosing from among the available options and executing an action associated with the selected option. The quality of the decision made by the collective depends upon the value of the chosen option.

Definition 5. Decision: A decision occurs when an acceptable majority of the agents complete an action, $act_i \in Act_{choose}$ associated with a single available option, ϕ_i, defined by $|A_{act_i}| \geq N - \delta$, where A_{act_i} is the set of agents performing act_i, and δ is the acceptable loss of agents during the decision problem.

There are five types of decisions. The most desirable outcome occurs when the collective executes the action associated with the highest valued option available. A collective that fails to find the optimal option is expected to select the best of the options that are discovered. Decisions are incorrect if the highest valued option found is not chosen. The collective fails if a decision is not made within a predefined time period. Finally, it is possible for collectives to split by selecting and executing two different options. Each decision is formally defined in Definitions 6 through 10.

Definition 6. Optimal Decision: An optimal decision arises when the implemented action is associated with ϕ_i and $i = \underset{j \in \{1,..,n\}}{\arg\max}\, v_j$.

Definition 7. Correct Decision: A correct decision is reached when the implemented action is associated with ϕ_i and $v_i = max_{\phi_j \in \Phi^{found}} v_j$, where $\phi_j \in \Phi^{found} \iff f_j = true$.

Definition 8. Incorrect Decision: An incorrect decision is reached when the implemented action is associated with ϕ_i and $v_i < max_{\phi_j \in \Phi^{found}} v_j$.

Definition 9. Failed Decision: A failed decision occurs when less than an acceptable majority executes an action associated with one of the available options ($|A_{act_j}| < N - \delta \; \forall j \in \{1,..,n\}$) after a maximum time limit.

Definition 10. Split Decision: A split decision represents a frustrated state during explicit collective action selection in which multiple actions are executed, defined by $\sum_{i \in \{1,..,n\}} |A_{act_i}| = N \bigwedge |A_{act_i}| \neq N \; \forall i \in \{1,..,n\}$.

III.2 Decision Making and Collective Action Selection Dynamics

The collective is a homogeneous group that pools information within a shared hub in a manner similar to decisions made in social insect colonies (e.g., Seeley, 2010; Mallon et al., 2001). Individual robots are cooperative, which means that there is negligible conflict of interest between them (Conradt and Roper, 2005). The values of different options are not directly compared and individuals are required to independently assess the values of the options before changing their opinion (Parker and Zhang, 2009; Reina et al., 2015b; Seeley et al., 2012). Indirect option comparisons serve two important purposes. First, option assessments are assumed to be error prone and direct comparisons propagate errors within the collective. Second, direct value comparisons have the potential to lead to indefinite frustrated states by reducing the ability of the collective to break symmetry when choosing between multiple equally valued options (see Chapter II.2.2). The presented solution causes symmetry breaking scenarios to happen more frequently than in the original model. A collective using the original model that must choose between two equally valued sites, ϕ_A and ϕ_B, in which $cost_A < cost_B$, will choose ϕ_A, since the cost difference reduces the problem difficulty. The presented model reduces environmental bias, which causes symmetry breaking decision problems to arise between equally valued options with disparate costs.

III.2.1 Collective Best-of-*n* Macroscopic Model

A strength of Reina *et al.*'s model is its linkage between the macroscopic model of the collective and the emergent behavior resulting from interactions between the individual robots. The macroscopic model is a deliberative process derived from a recent value-sensitive honeybee nest site selection strategy (Pais et al., 2013; Seeley et al., 2012). All robots in the collective are either part of the uncommitted population (Ψ_U), which does not favor any option, ϕ_0, or part of the favoring population (Ψ_{F_i}) that supports option ϕ_i. The agents periodically transition between interactive and latent states. Agents are interactive while in the decision making hub and communicate only with other interactive agents. Latent agents do not interact with other agents or travel outside the hub to either explore or evaluate options. Equations III.1 and III.2 describe these populations.

$$\dot{\Psi}_{F_i}^{delib} = R_{\gamma_i}\Psi_U + R_{\rho_i}\Psi_{F_i}\Psi_U - R_{\alpha_i}\Psi_{F_i} - \sum_{j\neq i} R_{i,\sigma_j}\Psi_{F_i}\Psi_{F_j}, \tag{III.1}$$

$$\Psi_U = 1 - \sum_{m\in\{1..n\}} \Psi_{F_m}. \tag{III.2}$$

Robots are initially in the uncommitted population (Ψ_U) and search the set of options, Φ. Population Ψ_{F_i} favors ϕ_i and increases as uncommitted robots discover ϕ_i, at rate R_{γ_i}, or as uncommitted robots are recruited by robots that favor option ϕ_i, at rate R_{ρ_i}. Ψ_{F_i} decreases as its robots abandon ϕ_i at rate R_{α_i}. Ψ_{F_i} also decreases when its members are inhibited by robots committed to $\phi_{j\neq i}$, at rate R_{σ_j}. The superscript $delib$ in the growth of the population, $\dot{\Psi}_{F_i}^{delib}$, indicates the change in the population of favoring agents due to this deliberative process.

The transition rates in Equation III.1 must be adjusted to account for the portions of the populations that are latent and interactive. Latent favoring and uncommitted robots are the only populations that transition states as a result of abandonment (R_{α_i}) or discovery (R_{γ_i}), respectively. Interactive favoring robots attempt to recruit interactive uncommitted robots (R_{ρ_i}) and inhibit interactive robots that favor $\phi_{j\neq i}$. Uncommitted and favoring agents become interactive at rates $R_U^I = \bar{\tau}^{-1}$ and $R_{F_i}^I = \tau_i^{-1}$, respectively, where $\bar{\tau}$ is the expected average round trip time between the decision making hub and all options $\phi_j \in \Phi^{found}$ and τ_i is the round trip time to a specific option. The collective is assumed to not possess any information about options $\phi_j \notin \Phi^{found}$, which means $\bar{\tau}$ is either an estimate or chosen by the designer to control agent interaction rates.

The option specific interaction rate is used for the favoring agents in order to demonstrate the impact of different option locations on the collective's decision making strategy. The original model assigns all robot interactive and latent transition rates to R_U^I and R^L, respectively, where R^L is equal to $9R_U^I$ in order to ensure that approximately ten percent of the population is interactive at a given time (Reina et al., 2015b). Increasing the distance between ϕ_i and the decision making hub; however, lowers the interaction rate for Ψ_{F_i} in relation to options that are closer to the hub. The ratios r^L_U, r^I_U, $r^L_{F_i}$ and $r^I_{F_i}$ determine the portions of the populations in the uncommitted latent (U^L), uncommitted interactive (U^I), favoring latent (F^L_i), and favoring interactive (F^I_i) populations, respectively. The ratios are determined by the following

$$r^L_U = \frac{R^L}{R^L + R^I_U};\ r^I_U = \frac{R^I_U}{R^L + R^I_U}, \tag{III.3}$$

$$r^L_{F_i} = \frac{R^L}{R^L + R^I_{F_i}};\ r^I_{F_i} = \frac{R^I_{F_i}}{R^L + R^I_{F_i}}. \tag{III.4}$$

The ratios in Equations III.3 and III.4 are used to limit the transition rates between the populations shown in Equation III.1. The resulting adjusted transition rates for Equation III.1 are determined as follows

$$R_{\rho_i} = r^I_{F_i} r^I_U \rho_i;\ R_{i,\sigma_j} = r^I_{F_i} r^I_{F_j} \sigma_j, \tag{III.5}$$

$$R_{\gamma_i} = r^L_U \gamma_i;\ R_{\alpha_i} = r^L_{F_i} \alpha_i. \tag{III.6}$$

Equation III.5 limits the recruitment and inhibition interactions to the interactive portions of the uncommitted and favoring populations. Equation III.6 limits discovery and abandonment to the latent portions of the population. As ϕ_i increases in distance from the decision making hub, less of its favoring population is available to interact with other populations; thus, making that option more vulnerable to option abandonment than the closer options. The primary transitions in the original model are defined by the four primary agent transitions: recruitment (ρ_i), inhibition (σ_i), discovery (γ_i), and abandonment (α_i). The values of each of these transition strengths are adopted from prior work (Reina et al., 2015b; Cody and Adams, 2017b) and are described by

$$\rho_i = \sigma_i = v_i;\ \alpha_i = \alpha v_i^{-1}, \tag{III.7}$$

$$\gamma_i = v_i g(d_i) = \frac{v_i \mu e^{\xi d_i}}{d_i}. \tag{III.8}$$

Equation III.7 shows that the recruitment (ρ_i) and inhibition (σ_i) strengths are equal to option value (v_i). Abandonment strength (α_i) is inversely proportional to v_i, where α is the minimum abandonment rate (set to 5% per traveled meter (Reina et al., 2015b)). Equation III.8 shows that option discovery (γ_i) decays with distance d_i from the collective's deliberation area, where μ and ξ are 0.058 and -0.29, respectively. The values of μ and ξ were determined by fitting a curve to discovery data obtained from running simulations of the agent's exploration behavior within the environment, which is similar to prior work (Reina et al., 2015b; Cody and Adams, 2017b) and is described in Chapter III.3.

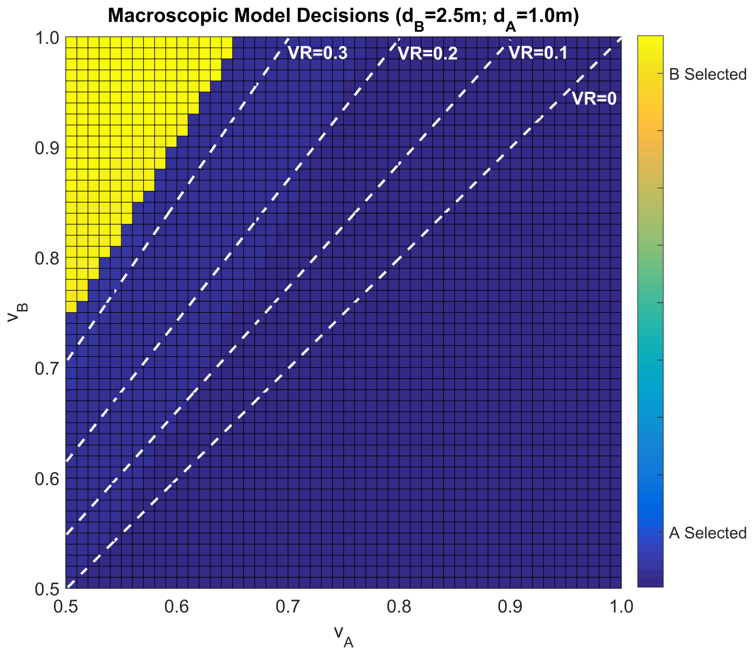

Figure III.1: The behavior of the original model is presented in a series of decisions in which ϕ_B has a greater value than ϕ_A, but a higher cost due to its greater distance. The dashed lines represent the decision boundaries for different value resolutions (VR).

The model behavior provided in Fig. III.1 demonstrates a series of decision problems between two options, A and B. The value of each option (v_A and v_B) are within the range $[0.5, 1.0]$, while the distances are fixed at $d_A = 1.0$m and $d_B = 2.5$m. Fig. III.1 demonstrates the effect of environmental bias, which is demonstrated using different option distances from the collective's decision making hub. The dashed lines indicate the decision boundary for different system value resolutions (Reina et al., 2015b).

Value resolution, VR, describes the sensitivity of the decision making model to both the difference between the values of the considered options and the magnitude of these values (Reina et al., 2015b; Pais et al., 2013). This property describes the minimum distinguishable difference between two options required for reliable collective selection of the higher valued option. Resolution is calculated using the following equation:

$$VR = \frac{|v_{min} - v_{max}|}{v_{max}}, \tag{III.9}$$

where the values v_{min} and v_{max} are set by the designer and are used to evaluate the system's ability to distinguish between options. A lower resolution value indicates the collective is increasingly able to choose reliably the higher valued option, even when the difference between the options is low.

Fig. III.1 demonstrates that ϕ_B is unlikely to be selected, even when v_B is much higher than v_A, due to

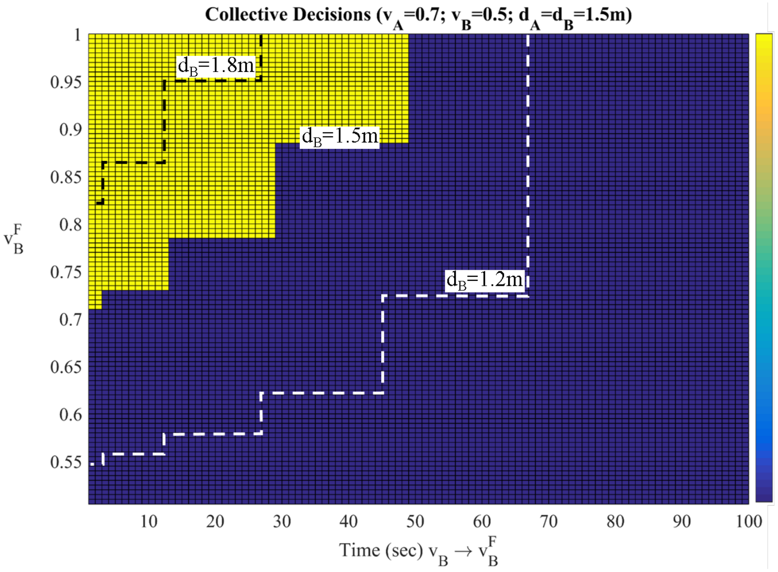

Figure III.2: The effect of environmental bias on the decision making process is shown under conditions of dynamic option value changes. The collective favors ϕ_A due to its lower cost, even though ϕ_B becomes the optimally valued option later in the decision making process.

ϕ_B's low discovery rate and higher round trip time. If the system only requires a resolution value greater than 0.3, the behavior shown in Fig. III.1 is considered acceptable. The macroscopic model assumes that the underlying system is infinite and well-mixed (Kampen, 2007; Reina et al., 2015b), which is unlikely to be true for robotic option selection trials, unless the individual agent behaviors have been carefully designed to encourage population mixing (Reina et al., 2016). Research questions R_1 and R_2 are addressed in Chapter III.2.2, which introduces model extensions that reduce environmental bias and attempt to move the decision boundary in Fig. III.1 to the $VR = 0$ line.

Negative environmental bias also impacts the collective best-of-*n* model's responsiveness to changes in the option value, as shown in Figure III.2. The collective chooses between options A and B, which have values $v_A = 0.7$ and $v_B = 0.5$, respectively, at the start of each decision problem. At a specific time, shown on the x-axis, v_B becomes v_B^F. Three boundaries are shown. When $d_A = d_B = 1.5$m, the collective is unable to select ϕ_B, if its value changes after 50 seconds (s). The collective is less likely to choose ϕ_B when it is located further away than ϕ_A, which is shown by the boundary at 1.8 meters (m). The collective is significantly more responsive when ϕ_B is located closer to the deliberation area than ϕ_A, which is shown by the boundary at 1.2m. Research question R_3 is examined in Chapter III.3.2, in which the ability of both the original collective best-of-*n* model and the model extensions described in Chapter III.2.2 are evaluated for their responsiveness to changes in the option value during the decision making process.

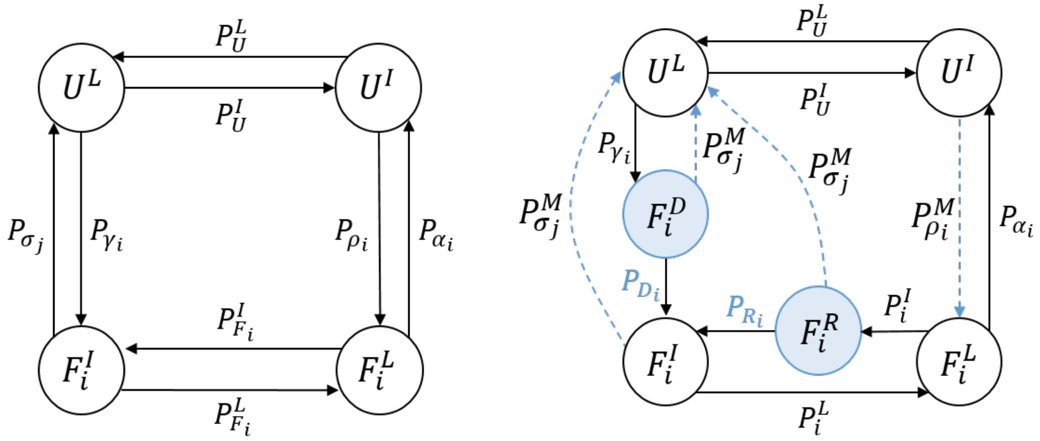

(a) Original Model (Reina et al., 2015b)

(b) Extended Model

Figure III.3: A comparison of Reina *et al.*'s (2015b) original individual agent model and the extended model. Extensions shown in Fig. (b) are indicated by shaded control states and dashed transitions.

III.2.2 Bias Reducing Model Extensions

Two extensions are presented to reduce environmental bias. The first extension, *interaction delay*, addresses research question R_1 by imposing semi-latent control states with exit probabilities that are determined by the option's distance to the collective's deliberation area. The second extension, *interaction frequency modulation*, addresses R_2 by increasing individual robot interaction attempts in relation to the location of the considered option. The extensions modulate robot interaction activities in order to reduce environmental bias by compensating for disparate discovery times and round trip travel times for options in the environment.

The original and extended models are presented in Figs. III.3a and III.3b, respectively. Shaded control states and dashed transitions indicate the extensions in Fig. III.3b. Robots in both models are uncommitted (U^*) or favoring (F_i^*), and are either interactive or latent, as indicated by the I and L superscripts, respectively. State transitions are stochastic, with probabilities indicated by P_*, that are equal to the rates transition rates described in Chapter III.2.1.

According to the original model, in Fig. III.3a, uncommitted, latent agents (U^L) explore the environment and transition to an interactive favoring state (F_i^I) after discovering ϕ_i with probability P_{γ_i} and returning to the decision making hub. Favoring interactive agents supporting ϕ_i recruit uncommitted interactive agents (U^I) with probability, P_{ρ_i}. Interactive agents in state F_i^I return to an uncommitted state when inhibited by agents in state $F_{j \neq i}^I$ with probability P_{σ_j}. Finally, latent favoring agents (F_i^L) travel from the decision making hub to ϕ_i in order to reassess v_i and potentially abandon ϕ_i with probability P_{α_i}. Agent interactions only occur just before agents transition from interactive to latent states in order to ensure that each agent's set of neighbors changes frequently between interactions (Reina et al., 2015b).

The interaction frequency modulation extension is indicated in Fig. III.3b using the superscript, M, in the inhibition ($P^M_{\sigma_j}$) and recruitment ($P^M_{\rho_i}$) transitions. The purpose of this extension is to address the discrepancy in transition rates between interactive and latent states for agents favoring options that are located at different distances from the collective's decision making hub (see Chapter III.2.1). Agents that favor options closer to the decision making hub interact more frequently than agents supporting options further from the hub. Interaction frequency modulation increases agent interaction attempts by $M_i = r(d_i/d_{max})$, where d_{max} is the maximum range of the collective's robots and r is chosen to be the ratio of the expected maximum round trip time (determined by d_{max}) to the minimum round trip time (determined by the dimensions of the decision making hub).

The second extension introduces the control states F^D_i and F^R_i in order to impose delays on favoring agents supporting options with lower discovery and round trip travel times, respectively. Uncommitted latent agents that discover ϕ_i enter a delaying control state (F^D_i), which they exit with probability P_{D_i}, in order to reduce the effect of disparate discovery times on the decision making process. Favoring latent agents that return to the nest transition to state F^R_i, which slows the agent's transition to the interactive state. Exit probabilities for the control states are determined by

$$P_{D_i} = \frac{\rho_i}{g(d_{max})^{-1} - g(d_i)^{-1}},$$

$$P_{R_i} = \frac{\rho_i}{\tau_{max} - \tau_i},$$

where τ_{max} is $2d_{max}$, or the maximum round trip time. Further, it is assumed that $d_i < d_{max} \forall i \in \{1..n\}$. Both states are semi-latent, which means robots in these states do not initiate either recruitment or inhibition interactions, but are vulnerable to inhibition by other robots. The experimental results presented in Chapters III.3.1 and III.3.2 indicate that these extensions reduce bias (R_1), improve decision efficiency (R_2), and enable uniform responsiveness to dynamic option value changes (R_3).

The presented solution to the issue of negative environmental bias assumes individual agents are able to travel at a constant velocity within the collective's search space. When this assumption is valid (e.g., in simple environments), the presented solution directly addresses the decision bias induced by disparate option distances to the collective's decision making hub. The approach does not; however, address several realistic sources of environmental bias that are not homogeneous throughout the search space. Obstacles and structures are expected to impose uneven travel delays that depend on the option location and the path the individual agent took to reach the option. Rough terrain, slope, and vegetation are additional sources of environmental bias that vary across a realistic search space and are unlikely to be known to the entire collective. Finally, several sources of environmental bias are likely to change over time. Wind speed, current, and collective congestion each will bias collective decisions in more difficult scenarios. The presented solution approach assumes the collective is limited to a circular search space that is centered on the collective decision making hub, with a radius that represents the maximum distance between the hub and an acceptable option. Individual agents compare the distance between their current option and the collective decision making hub to the maximum distance and alter their interaction rates in order to reduce the effect of the option's distance on the collective decision making process. The approach is readily modified to accommodate more difficult terrain, where agents can adjust their interaction behavior based on their actual, rather than ideal, travel time to their supported option. Comparison of the agent's travel time to a maximum travel time parameter enables similar functionality without relying on a homogeneous search space.

III.2.3 Explicit Collective Action Selection

Explicit action selection requires the collective to recognize when it has made a best-of-*n* decision and transition to task implementation. After which, the collective must return to an information gathering state in order to choose and implement its next action. The presented approach extends the collective best-of-*n* decision strategy from Chapters III.2.2 and III.2.2 with Parker and Zhang's (2010) task sequencing framework (see Chapter II.2.2). The original task sequencing framework evaluated a collective conducting area clearing activities. The robots were required to determine when the task was complete and transition to the next task as a group. A voting based quorum detection mechanism was used by individual robots to estimate the collective's task completion. A robot that sensed the task was completed began querying neighbors to determine if its neighbors agreed that the collective's task was complete. Each robot maintained a queue of voting messages and entered a committed state, once a threshold of agreeing messages in the queue was received. Once in the committed state, robots sent commitment messages to neighbors. Robots that received a commitment message responded with an acknowledgment for the first message and ignored subsequent messages. Finally, robots initiated timers on receipt of an acknowledgment and transitioned to the next task if the timer expired (Parker and Zhang, 2010).

The presented approach introduces three important changes to the task sequencing framework. First, the collective uses a best-of-*n* strategy to choose each action prior to execution, rather than knowing the sequence in advance. Second, the collective is dispersed in a wide area and only communicates within the decision making hub. Unlike in Parker and Zhang's 2010 research, the collective risks making a decision while some robots are outside the hub. When the chosen task requires the collective to move, robots that are outside the decision making hub when a decision is reached risk being left behind. Finally, the presented approach is intended to be used in large (> 50) groups of agents, which increases the likelihood of quorum estimation noise resulting from duplicate messages in agent queues and from fluctuations in the representation of options

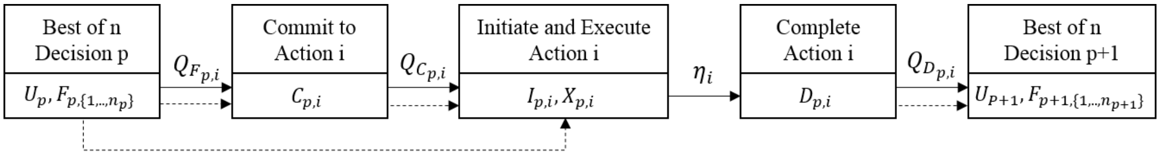

Figure III.4: A high level overview of the presented approach to explicit collective action selection. The collective chooses and executes actions in sequence. Solid lines indicate spontaneous transitions, while dashed lines indicate transitions that require the receiving population to interact with the losing population.

within a given agent's neighborhood of peer agents, when it attempts to interact (see Chapter II.2.2).

A visual representation of the presented approach is provided in Figure III.4, which shows the collective deliberating between actions, choosing an action, executing and completing the action, and finally transitioning to the next decision. Solid lines represent spontaneous transitions between states, while dashed lines represent transitions that require the receiving population to interact with the losing population. The interaction between the populations is described by the following system of ordinary differential equations:

$$\dot{\Psi}_{F_{p,i}} = \dot{\Psi}^{delib}_{F_{p,i}} - r^I_{F_i}\Psi_{F_{p,i}}\left[\omega_{F_{p,i}} + \sum_{m\in\{1,..,n\}}\left(\Psi_{C_{p,m}} + \Psi_{I_{p,m}}\right)\right], \tag{III.10}$$

$$\dot{\Psi}_{C_{p,i}} = \omega_{F_{p,i}} r^I_{F_i}\Psi_{F_{p,i}} + \Psi_{C_{p,i}}\left[r^I_U\Psi_{U_p} - \omega_{C_{p,i}} + \sum_{m\in\{1,..,n\}}\left(r^I_{F_m}\Psi_{F_{p,m}} - \Psi_{I_{p,m}}\right)\right], \tag{III.11}$$

$$\dot{\Psi}_{I_{p,i}} = \omega_{C_{p,i}}\Psi_{C_{p,i}} + \Psi_{I_{p,i}}\left[r^I_U\Psi_{U_p} - \omega_{I_{p,i}} - \frac{r^I_U}{kT} + \sum_{m\in\{1,..,n\}}\left(r^I_{F_m}\Psi_{F_{p,m}} + \Psi_{C_{p,i}}\right)\right], \tag{III.12}$$

$$\dot{\Psi}_{X_{p,i}} = \Psi_{I_{p,i}}\left(\omega_{I_{p,i}} + \frac{r^I_U}{kT}\right) - \eta_{p,i}\Psi_{X_{p,i}} \tag{III.13}$$

$$\dot{\Psi}_{D_{p,i}} = \eta_{p,i}\Psi_{X_{p,i}} - \Psi_{D_{p,i}}\left(\omega_{D_{p,i}} + r^I_U\Psi_{U_{p+1}}\right) \tag{III.14}$$

$$\dot{\Psi}_{U_p} = \Psi_{D_{p-1,i}}\left(\omega_{D_{p-1,i}} + r^I_U\Psi_{U_p}\right) - \dot{\Psi}^{delib}_{F_{p,i}} - r^I_U\Psi_{U_p}\left(\Psi_{C_{p,i}} + \Psi_{I_{p,i}}\right). \tag{III.15}$$

The collective is a population of agents, $\Psi = 1$, which is initially uncommitted, $\Psi_{U_p} = 1$, and initiates a collective action selection problem, CAS_p, where $p \in \mathbb{N}$ identifies a specific decision. Equation III.10 represents that the population favoring $\phi_i \in \Phi^{found}$ increases according to the deliberative process, $\dot{\Psi}^{delib}_{F_{p,i}}$ (see Equation III.1), where n_p is the number of available options for CAS_p. All interactions with the favoring and uncommitted population are limited to the interactive portions of each population, as described by the ratios r^I_U and $r^I_{F_i}$, respectively (see Equation III.4). Population Ψ_{F_i} senses a quorum supporting action i at rate $r^I_{F_i}\omega_{F_{p,i}}$ and transitions to the committed population, $\Psi_{C_{p,i}}$. Equation III.11 allows $\Psi_{C_{p,i}}$ to increase by committing the uncommitted and favoring populations to action i through interactions. The committed population transitions to initiation, $\Psi_{I_{p,i}}$, after sensing a quorum of committed agents, at rate $\omega_{C_{p,i}}$, which ensures that enough of the population has stopped deliberating and is ready to initiate a task. The initiating population, as shown in Equation III.12, increases through interactions with the committed and deliberating populations in order to transition those populations to the initiating state. The initiating population begins execution, $\Psi_{X_{p,i}}$, after either sensing a quorum of initiating agents, at rate $\omega_{I_{p,i}}$, or after a timeout period, r_U^I/kT, where k is the length of the message queue agents use to test that a quorum exists and T is the quorum threshold (not shown in Figure III.4). Initiating agents that do not receive any messages in a period of time defined by ${r_U^I}^{-1}$, add their own vote to their voting queue. After at most kT periods the individual agent transitions. The executing population $\Psi_{X_{p,i}}$ enters the completed, or "Done", population, $\Psi_{D_{p,i}}$ at the rate of action completion, $\eta_{p,i}$. Finally, the completed population transitions to the uncommitted state for the next action selection problem, CAS_{p+1}, upon sensing a quorum at rate $\omega_{D_{p,i}}$ and through interaction with the uncommitted population $\Psi_{U_{p+1}}$.

Equations III.14 and III.15 assume agents ignore messages from previous decisions, which is a simplifying assumption similar to that in prior work (Parker and Zhang, 2010). The reason for this assumption is to prevent uncommitted agents in the next decision process $\Psi_{U_{p+1}}$ from interacting with populations in earlier decision states (e.g., $\Psi_{I_{p,i}}$) and initiating the previous selected action a second time. Agents in the $\Psi_{U_{p+1}}$ population are likely to be in the vicinity of agents in population $\Phi_{I_{p,i}}$ in two different scenarios. The first scenario occurs when the collective decides between actions that are executed within or near the collective's current decision making hub. The second scenario occurs when the collective moves its decision making hub after every decision and chooses to occupy a previous location a second time.

The quorum sensing rate, $\omega_s \ \forall \ s \in \{F_{p,i}, C_{p,i}, I_{p,i}, D_{p,i}\}$, depends upon the applied quorum detection rule. The following set of equations determine the parameter values:

$$\omega_s^M = \sum_{j=\lceil kT \rceil}^{k} \begin{pmatrix} k \\ j \end{pmatrix} \Psi_s^j (1-\Psi_s)^{k-j}, \tag{III.16}$$

$$\omega_s^K = \Psi_s^k. \tag{III.17}$$

The quorum sensing rate, ω_s, is determined by one of two functions, each of which depend upon the size of the applicable population that is sensing for the quorum, Ψ_s, and the length of the queue of messages, $k = 15$, retained by the agents (Parker and Zhang, 2009). A majority queue rule, ω_s^M, represents the probability that at least $\lceil kT \rceil$ votes in the queue match s, where T is the threshold value equal to 0.75 (Parker and Zhang, 2009). The k-unanimity rule requires all k votes to match s, which is equivalent to the majority rule when $T = 1$ (Scheidler et al., 2016).

The collective behavior for a set of single decision scenarios using the system described by Equations III.10 through III.15 is presented in Fig. III.5. The collective is predicted to unanimously implement an

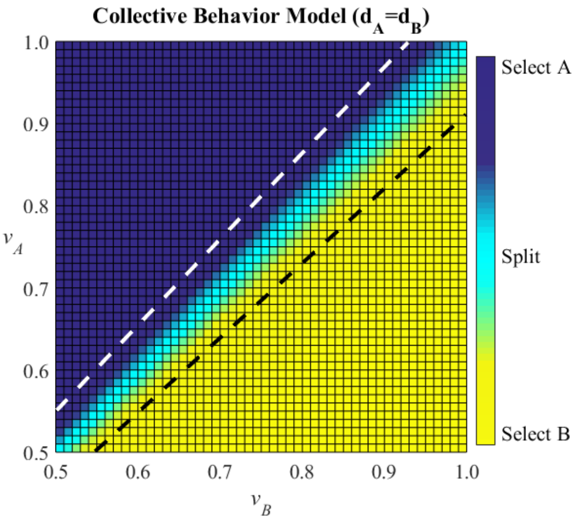

Figure III.5: The proposed model predicts that when ϕ_A and ϕ_B are equidistant, the collective will select the higher valued option (e.g., select ϕ_A when $v_A > v_B$). Split decisions are more likely along the $v_A = v_B$ boundary for the k-unanimity rule and between the dashed lines for the majority rule.

action associated with the higher quality option, ϕ_B, in most scenarios, but it is possible for the collective to split between competing options. Split decisions are more likely in the light blue area along the line $v_A = v_B$ for the k-unanimity rule and between the dashed lines for the majority rule. Although the stochastic nature of the underlying population is sufficient to enable a majority favoring one option, even when options are equally valued (Reina et al., 2015b; Pais et al., 2013), dual commitment is possible given the probability that agents will sense a false quorum (Parker and Zhang, 2009). Quorum threshold controls the speed and accuracy of the decision making process (Parker and Zhang, 2009). While Figure III.5 indicates that the stricter quorum sensing mechanism improves accuracy, it is expected to cause prohibitively long decision times for the collective. An examination of quorum detection techniques using the majority and k-unanimity thresholds as well as different voting message frequencies, described in Chapter III.2.3.1, and are examined in an experiment described in Chapter III.3.3.

III.2.3.1 Explicit Action Selection Individual Model

The individual behavior model for explicit collective action selection is an extension of the collective best-of-n process described in Chapters III.2.1 and III.2.2. Individual agents exist in one of six states during a decision problem, which correspond to the populations described in Chapter III.2.3: U_p, $F_{p,i}$, $C_{p,i}$, $I_{p,i}$, $X_{p,i}$, $D_{p,i}$. Quorum detection is necessary for initial agent transitions to the committed ($C_{p,i}$), initiating ($I_{p,i}$), executing ($X_{p,i}$), and completed (D_i) states. Individual agents maintain a queue of the last k messages received. A quorum is detected when the ratio of messages that match the agent's state to the total messages in the queue exceed a given threshold, T.

Queue message length, like quorum thresholds, control the speed and accuracy of the decision making process (Campo et al., 2011; Parker and Zhang, 2011). The primary difference between the parameters is that as queue lengths increase, agent quorum detection is influenced by older information (Parker and Zhang, 2010). The presented approach considers two different types of queues, persistent and episodic, in order to investigate this phenomenon (see Chapter III.3.3). Persistent queues are maintained by agents throughout the decision problem, without regard to an agent's change in state. Episodic queues are cleared each time an agent transitions to a new state. Due to these frequent state changes, voting messages must be received more often than regular interaction messages in order to provide agents an opportunity to sense a quorum prior to changing state. Voting messages are sent by favoring interactive agents with a frequency of rP_L, where $r > 1$.

The individual behavior model is shown in Figure III.6. Agents in the uncommitted and favoring states behave according to the collective best-of-*n* strategy, described in Chapter III.2.1 and are only interactive within the decision making hub. A favoring agent ($F_{p,i}$) spontaneously transitions to the committed state ($C_{p,i}$) after detecting a quorum of support for ϕ_i with probability $P_{\omega_{F_{p,i}}}$. Committed agents remain within the nest area and they commit deliberating agents that they encounter. Committed agents spontaneously transition to the initiating state, $I_{p,i}$, with probability $P_{\omega_{C_{p,i}}}$. The initiating population interacts with agents in the previous three states, causing them to enter the initiation state. Agents transition to the execution state, $X_{p,i}$ after detecting a quorum, with probability $P_{\omega_{I_{p,i}}}$. Agents in state $I_{p,i}$ add their own vote to their queue with probability P_U^I if no messages have been received, which ensures that most initiating agents start executing action *i* with a probability of P_U^I/kT. This mechanism is particularly important when the chosen task requires the collective to change the location of its decision making hub. Without resorting to a special broadcast message, the synchronous initiation of action execution is not guaranteed and a small number of agents are expected to be left in the original hub when the rest of the agents move to the new location. The self voting mechanism used by agents in state $I_{p,i}$, prevent initiating agents from remaining in the original location for extended periods of time, while they attempt to detect a quorum. Agents complete act_i with probability $P_{\eta_{p,i}}$ and spontaneously transition to the completed state, $D_{p,i}$. Finally, once completed agents detect a quorum, they transition to the next decision state's uncommitted population, U_{p+1}.

The presented behavior model does not solve the problem of duplicate quorum detection messages. Several existing collective decision making approaches assume anonymous messaging between individual robots, which is intended to reduce the complexity of individual agents and avoid requiring that individuals have unique identifiers (Parker and Zhang, 2009, 2010; Reina et al., 2015b). The use of autonomous messaging during quorum detection implies that individual robots are unable to distinguish between duplicate and original messages received from neighboring robots. Duplicate messages during quorum sensing reduces the accuracy of each individual's estimate of the collective state. Although it was not evaluated in their research,

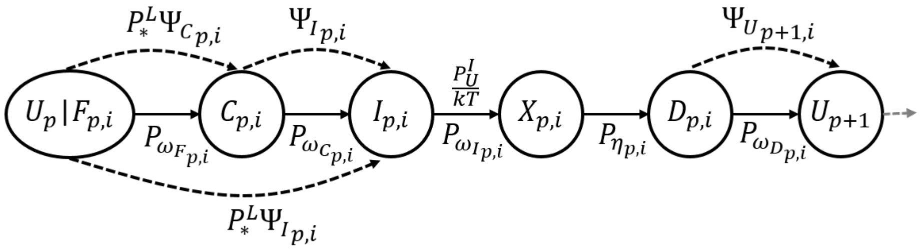

Figure III.6: The individual behavior model for the presented approach to explicit collective action selection. Solid lines are spontaneous transitions between states, while dashed lines require interactions between agents in the receiving and losing states. The term P_*^L is used for convenience to represent the latent transition probabilities for both U_p and $F_{p,i}$.

Parker and Zhang (2010) recommended that agents delay between interactions in order to reduce the chances of receiving duplicate messages. The drawback to this approach is that extending time between interactions increases the collective's decision making time. Another solution includes additional information in voting messages, such as sending agent identification in order to prevent or limit duplicate messages (Hamar and Dove, 2012).

Three different approaches to the problem of duplicate quorum voting messages are considered and evaluated. First, individuals perform no message filtering, but increase the period between interactions in order to reduce the likelihood of receiving duplicate messages. Second, individuals use weak voting message filtering and generate a random number in the range $[0, k-1]$ each time they enter the interactive state and include this number in their voting messages. Agents receiving this message ignore messages with values already in the queue during the current interactive period. Finally, individuals use strong voting message filtering and generate a random interaction identifier each time they enter the interactive state and duplicates are ignored. The effectiveness of the individual behavior model in general, and the effectiveness of each of these approaches is the purpose of the experiment described in Chapter III.3.4, which addresses the fourth, and final, research question for independent explicit collective action selection in robotic collectives.

III.3 Simulated Robotic Collective Experiments

Four evaluations were conducted in order to analyze the effectiveness of the solutions for addressing research questions R_1-R_4. The experiments consisted of unique models, metrics, and environments that are described in Chapters III.3.1 through III.3.4. Robot models were implemented in Java in order to permit interaction with a variety of simulation environments. Processing, an open-source programming environment, was used for all experiments (Reas and Fry, 2014).

The options considered in all experiments were sites (see Definition 3) located within a two dimensional environment. Both sites and agents in these experiments were represented as point masses. The robot agent

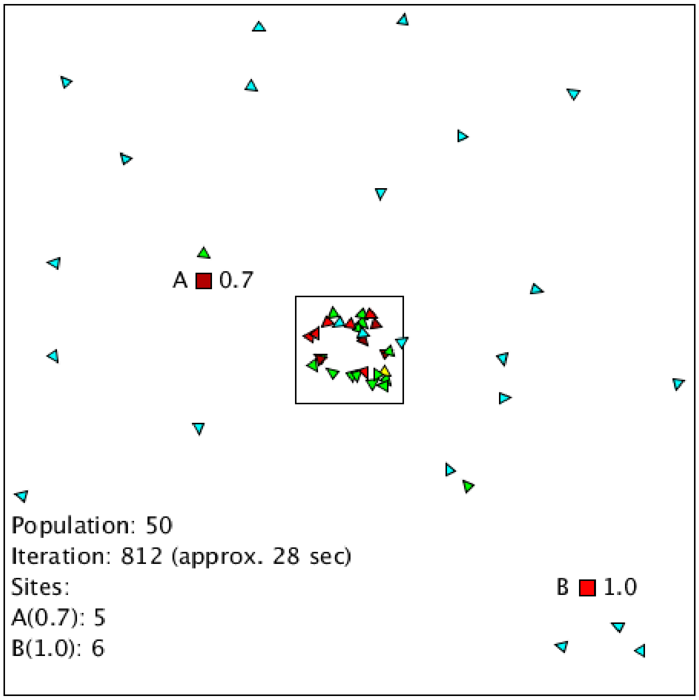

Figure III.7: The two dimensional, point-mass simulation environment during a decision problem in which the optimal site, *B*, is further from the decision making hub than site *A*. The central box is the decision making hub where agents interact.

characteristics were based on the e-puck robot, using local communication through infrared sensors (Mondada et al., 2009). The simulated robots have a sensory range and communication range of 12.8cm, a constant speed of 10cm/s, and move through the environment using a correlated random walk, when not traveling between the nest area and a known site (Reina et al., 2015a; Codling et al., 2008). Agents travel directly to and from the nest area and known locations in straight lines using perfect odometry, as in prior work (Reina et al., 2015a). Agents perform instantaneous site evaluations when they are within range of a site.

III.3.1 Experimental Analysis of Models with Static Site Values

The purpose of the static site values experiment was to determine whether the interaction frequency modulation and interaction delay extensions to the original model (see Chapter III.2.2) address research questions R_1 and R_2. Two hypotheses were examined. The first hypothesis, H_1, was that interaction delays improve collective decision accuracy over the original model. The second hypothesis, H_2, was that interaction frequency modulation improves decision time and accuracy in the presence of environmental bias.

Table III.1: Evaluated Models

Model	Description
M_0	Reina et al.'s (2015b) Original model
M_1	M_0 with interaction frequency modulation
M_2	M_0 with interaction delay
M_3	M_1 and M_2 combined

III.3.1.1 Experimental Design

The simulation environment, shown in Fig. III.7, consisted of a 3.5×3.5m two dimensional, obstacle free search space with a central 0.7×0.7m nest area (where 1 pixel = 1cm). Agents were allowed to pass through the boundaries in order to prevent wall boundary interactions from improving site discovery probabilities. All probabilities were adjusted to account for five simulation iterations per second. The random walk behavior and the simulation environment affected the discovery rate parameters, μ and ξ, as previously discussed (see Equation III.7 in Chapter III.2.1). These parameter values were determined by fitting a curve to discovery time data for 1000 simulations in which agents used the correlated random walk to explore the environment for sites located at different distances from the decision making hub ($[0.5, 3.5]$m, at 0.5m increments).

The primary independent variable was the evaluated model, as summarized in Table III.1. The original model was Reina et al.'s (2015b) model. Model M_1 extended M_0 with interaction frequency modulation. Model M_2 extended M_0 with interaction delays. Finally, Model M_3 combining both extensions with M_0. The secondary independent variables included the collective population size, $N \in \{50, 100, 500\}$, the distance between the nest area and each of two sites, $d_i \in [0.5, 3.5]$m, in 0.5m increments, and the value of each site, $v_i \in \{0.6, 1.0\}$, incremented by 0.1. Each primary independent variable was paired with each combination of secondary independent variables and evaluated via 25 trials. Each trial was a decision problem in which the agents started within the nest area and explored the arena for two sites. The sites were detected when an agent was within sensory range.

The metrics for evaluating the performance of these extensions were success rate, decision time, and effectivity (Reina et al., 2015b). Success rate represented the percentage of trials for each combination of the secondary independent variables in which the collective achieved a quorum of support ($T = 0.75$) for the highest valued site. When sites were equally valued, an accurate decision was one in which either site was selected. Decision time was the length of time required to achieve a quorum for a site within the arena, whether the decision was accurate or not. Finally, effectivity was the percentage of total trials in which a quorum was reached prior to a maximum time limit. The imposed time limit was $12,000$ seconds (approximately 3.3hrs).

Table III.2: Success Rate (*SR*) and Decision Time (*DT*)

Model	Statistic	SR (%)	DT (min)
M_0	Mean (SD)	0.72 (0.42)	15.2 (23.0)
	Median	1.0	7.67
	Min (Max)	0 (1.0)	1.81 (199.3)
M_1	Mean (SD)	0.711 (0.407)	10.5(13.33)
	Median	1.0	7.00
	Min (Max)	0 (1.0)	1.98 (199.2)
M_2	Mean (SD)	0.877 (0.24)	25.38 (25.9)
	Median	1.0	17.0
	Min (Max)	0 (1.0)	4.46 (199.9)
M_3	Mean (SD)	0.875 (0.20)	16.5(13.4)
	Median	0.96	13.42
	Min (Max)	0.0 (1.0)	4.31 (199.3)

III.3.1.2 Results

The descriptive statistics for the success rates and decision times achieved by the models are provided in Table III.2. M_2 made the most accurate and slowest decisions, while M_1 was the fasted and least accurate. The interaction frequency modulation mechanism in M_3 improved its decision time, while maintaining a success rate comparable to M_2. The minimum (Min) and maximum (Max) success rates, in Table III.2 demonstrate that each model failed (0) or succeeded (1.0) every trial for at least one problem scenario. Trials that resulted in failed decisions, meaning that the collective did not reach a 75% majority within the imposed time limit, represented 3.5% of the total trials and were omitted from Table III.2. Model M_2 achieved an effectivity of 0.93, which means that it did not make a decision in 7% of its total trials. M_0, M_3, and M_1 achieved effectivity rates of 0.96, 0.98, and 0.99, respectively.

A Kruskal-Wallis test indicated a significant effect across the four models for success rate, $\chi^2(3, N = 157500) = 4534.6, p < 0.001$, and decision time, $\chi^2(3, N = 157500) = 29739, p < 0.001$. Significant effects were also observed between all models in pairwise comparisons using a Tukey and Kramer (Nemenyi) test for success rate ($p < 0.001$) and decision time ($p < 0.001$).

Site location, value, and population influenced the success rates of all models. The effective resolution boundaries by model and collective size are provided in Fig. III.8. Lower effective resolution boundaries indicate higher success rates. Increasing the differences in site distances from the decision making hub (x-axis), for all population sizes significantly reduced the success rates of M_0 and M_1, which were unable to select the optimal site when sites differed in distance by 1.5m or more. M_2 and M_3 were more resilient to differences in site locations. Increases in population size improved M_2's success rate the most, but had little impact on the success rates of the remaining models.

The skewed distributions for model decision times, as shown in Fig. III.9, demonstrate the challenge of

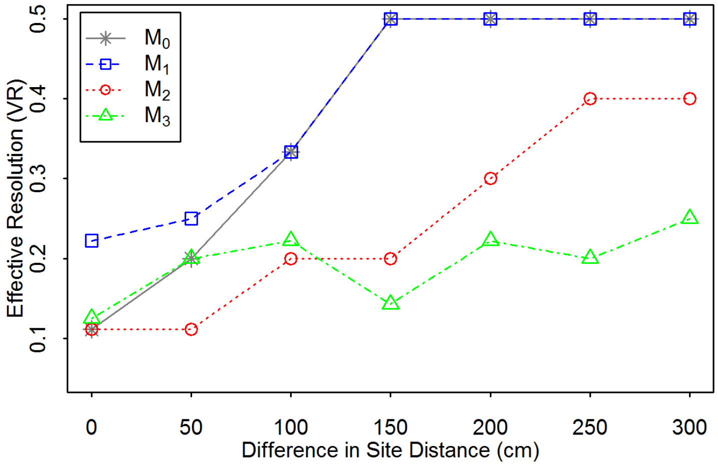

(a) Effective resolution (50 agents).

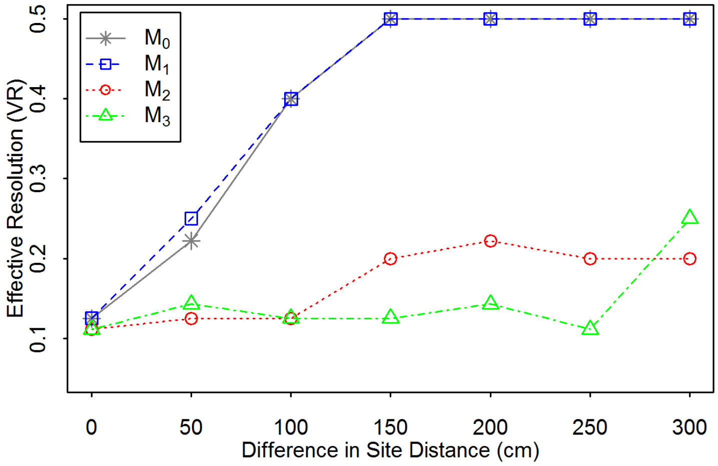

(b) Effective resolution (100 agents).

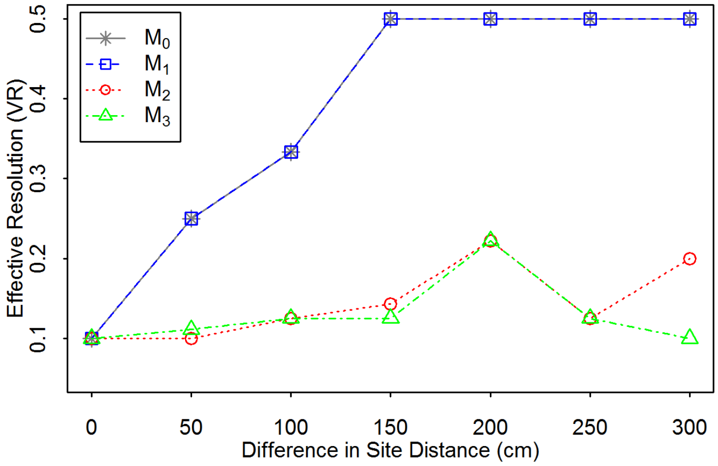

(c) Effective resolution (500 agents).

Figure III.8: The four models compared by success rate according to population size and differences between site distances from the decision making hub. Each model's effective resolution boundary is presented, in which lower lines indicate higher success rates at different site distances.

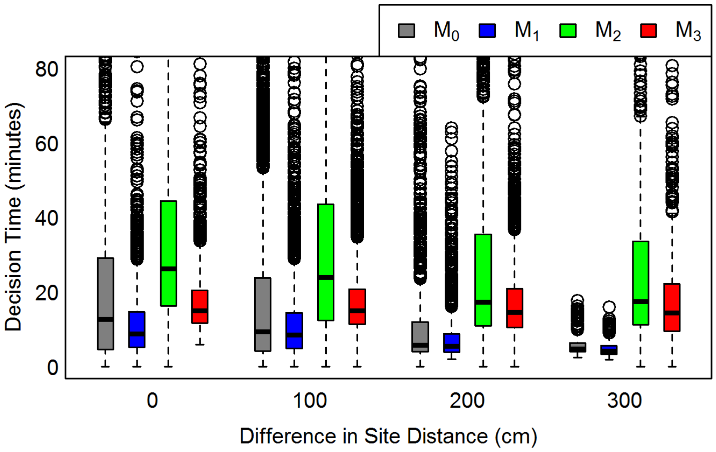

(a) Decision time for 50 agents (Outliers > 80min: M_0-638, M_1-88, M_2-1334, M_3-168).

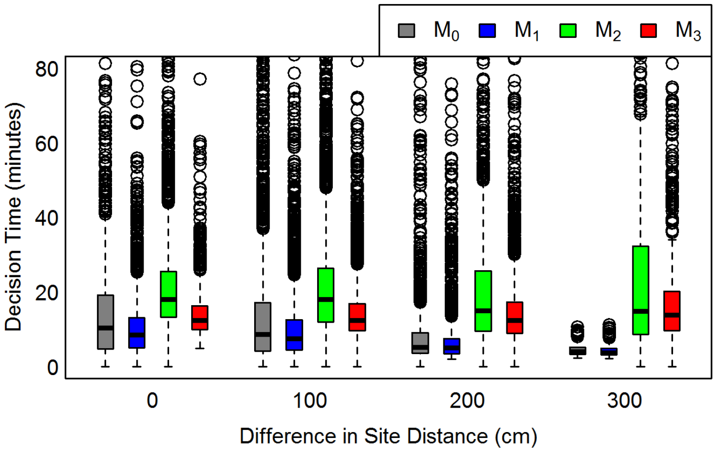

(b) Decision time for 100 agents (Outliers > 80min: M_0-332, M_1-70, M_2-390, M_3-89).

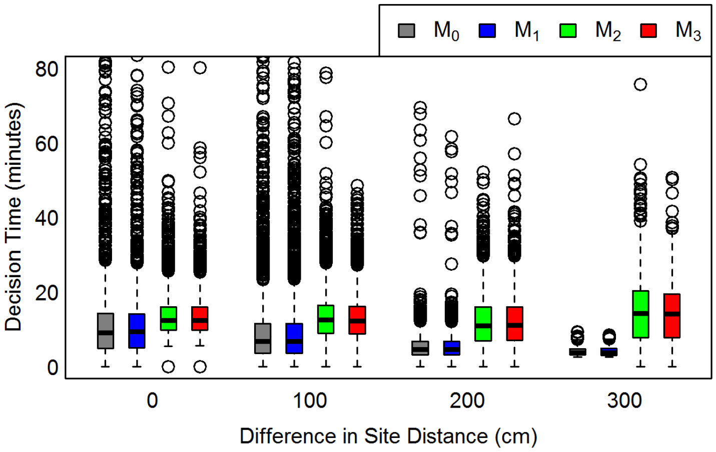

(c) Decision time for 500 agents (Outliers > 80min: M_0-142, M_1-101, M_2-52, M_3-56).

Figure III.9: The four models compared by decision times according to population size and differences between site distances from the decision making hub. The slowest 3,460 (2.1% of total trials) trials are omitted due to Decision Times that exceeded 80min in order to facilitate viewing the primary results.

predicting convergence time during collective site selection. The figures portray trials that completed prior to the imposed time limit. M_1 made the fastest decisions for all population sizes. Increasing difference in distance between the sites and the decision making hub improved the decision time of both M_0 and M_1, as each model rapidly converged to the closer site despite that site's value. M_2 and M_3 decision times were resilient to differences in site distances from the decision making hub. Increasing the collective population size resulted in a modest decrease in decision times for all models.

III.3.1.3 Discussion

The results indicate that incorporating interaction delay, in models M_2 and M_3, reduced environmental bias caused by distance for value-sensitive collective site selection problems. The results support hypotheses H_1 and partially supported H_2, but the evaluation of each individual extension deviated slightly from expectations. Interaction delays, as anticipated, improved success rate, at the cost of decision time and decreased effectivity. Interaction frequency modulation improved decision time as expected, but did not significantly affect success rate. The design decision to limit interaction frequency via the latency probability reduced the relative influence of site location for each population's frequency calculations. Additionally, the models implementing interaction delays (M_2 and M_3) experienced fluctuations related to spatial differences between sites. The interaction delay extension included discovery probability estimates based on a population size of 50 agents, which biased the models in favor of sites that were further away in several scenarios with greater populations. This bias was more noticeable in the combined model (M_3) and caused a slight reduction in its success rates overall.

While the presented results must be supported with future robotic experiments, their implications to robotic collective decision making are significant. First, combining different behavior modulation techniques is an effective means of improving accuracy without incurring costly decision time penalties. Second, value resolution is at least partially dependent on the environmental bias (e.g., distance). The combined model reduced this influence in an ideal setting, but future systems must be evaluated throughout the search range to provide reasonable accuracy guarantees. Third, although M_3, the combined model, achieved comparable overall decision speed to the original model, M_3 is expected to face more symmetry breaking scenarios and perform more slowly than the original model when the closer site is superior to the further site. Due to the combined model's relative indifference to site placement, it faces a symmetry breaking scenario any time two equal sites are within range, resulting in a decision time greater than the original model (Reina et al., 2016; Hamann et al., 2012). The original model is more susceptible to environmental bias and only faces a symmetry breaking scenario when its considered sites are equidistant to the decision making hub and equally valued.

The static site values experiment demonstrated that the interaction delay and interaction frequency modulation extensions provide partial solutions to research questions R_1 and R_2. Interaction delays partially addressed R_1 by reducing the influence of discovery and travel cost associated with sites in the collective's search space, which enabled models M_2 and M_3 to make decisions based primarily on site value in more scenarios than in the other two models. Two important points are necessary. First, interaction delays only address environmental factors that influence the time associated with site discovery and travel time between the decision making hub and the site. Other costs, such as additional energy requirements (e.g., to travel through difficult terrain) do not necessarily affect either of the evaluated costs and are; therefore, not addressed by interaction delays. Second, this experiment, and the rest of the experiments in this dissertation, do not consider differing assessment times between sites, which has been discussed in prior studies (e.g., Valentini et al., 2014). While adding assessment cost to the interaction delay extension is expected to be straightforward, it was not evaluated or incorporated into the model. Interaction frequency modulation partially addressed R_2 by improving the decision times with only modest effects on the decision accuracy of both M_1 and M_3, when compared to M_0 and M_2, respectively. The benefit gained by interaction frequency modulation; however, was observed to decrease as the population of agents increased.

Static site values are common in studies examining opinion-based collective best-of-*n* approaches, but do not address the effect of changing site values during the decision making process. Predictable responses to changes in site values are particularly important for explicit collective action selection, since the collective risks executing the wrong action if it does not react to a sudden value change. Option values are likely to change as a result of changes in the environment and changes in priority when the collective is operating as part of a human-collective team. The dynamic site values experiment examines the same set of models for their responsiveness to changing site values.

III.3.2 Experimental Analysis of Models with Dynamic Site Values

The most difficult problems faced by the original model in the static site value experiment (see Chapter III.3.1) were problems in which the optimal site was located further from the decision making hub than a closer site with a lower value. The dynamic site values experiment was designed to address research question R_3 by evaluating the original and extended collective best-of-*n* models in scenarios that are subject to site value change during the decision making process. The hypothesis, H_3, was that the decision making process will be more uniformly responsive to site value changes when the effect of environmental bias is reduced.

III.3.2.1 Experimental Design

The experimental design was based on the same simulation arena, evaluated decision making models, and individual agent behavior parameters as was used in the static site values experiment (Chapter III.3.1.1). Sites had an associated initial value, v_i, and a final value v_i^F that they assume at time t^F in order to simulate a drastic change in site value during the decision making process. The decision problems were designed to require the collective to make decisions in the presence of environmental bias and site value changes.

The primary independent variable was the evaluated model. The secondary independent variables were the agent population size ($N \in \{100, 300, 500\}$), site initial quality ($v_i \in [0.5, 1.0]$, incremented by 0.1), site final quality ($v_i^F \in [0.5, 1.0]$), site distances ($d_A = 1.5$m and $d_B \in \{1.0, 1.5, 2.0\}$m), and quality change time ($t^F \in [0, 200]$s incremented by 25s). Scenarios were defined by the cost, in terms of distance, of the optimal site (d_{opt}) and the mediocre (d_{med}) site. Low, moderate, and high cost scenarios were those in which $d_{opt} < d_{med}$, $d_{opt} = d_{med}$, and $d_{opt} > d_{med}$, respectively. Each pairing of secondary and primary independent variables were evaluated for 50 site selection decisions. Three metrics were considered: success rate - the percentage of trials in which the collective chose the best site according to the site's final values (v_i^F), response time - the period of time between t^F and a correct decision, and total decision time.

III.3.2.2 Results

The descriptive statistics for Success Rate (*SR*), Decision Time (*DT*), and Response Time (*RT*) are provided in Table III.3. M_0 performed as expected and achieved slightly lower success rates than in prior work by Reina et al. (2015b) due to the introduction of dynamic value changes. M_1 was observed to have the fastest decision and response times of all the models, and achieved a similar success rate to M_0. M_2 made the slowest, most accurate decisions. M_3 achieved a similar success rate to M_2 with decision times that were within one

Table III.3: Descriptive Statistics by Model and Metric

Model	Statistic	SR	DT	RT
M_0	Mean (SD)	0.76 (0.42)	8.91 (4.11)	7.05 (4.0)
	Median	1.0	8.2	6.5
	Min (Max)	0 (1.0)	3.05 (50.53)	0.11 (44.97)
M_1	Mean (SD)	0.76 (0.43)	7.99 (3.59)	6.14 (3.5)
	Median	1.0	7.35	5.71
	Min (Max)	0 (1.0)	2.81 (51.07)	0.03 (49.4)
M_2	Mean (SD)	0.93 (0.26)	10.94 (4.36)	9.3 (4.2)
	Median	1.0	10.01	8.66
	Min (Max)	0 (1.0)	4.4 (60.8)	2.0 (58.3)
M_3	Mean (SD)	0.90 (0.3)	9.26 (3.1)	7.6 (3.02)
	Median	1.0	8.54	7.17
	Min (Max)	0.0 (1.0)	3.99 (39.15)	1.21 (36.23)

minute of M_0's decision times. Interaction frequency modulation for M_1 and M_3 improved decision times for both models without significantly affecting success rate, as compared to M_0 and M_2, respectively. Interaction delays increased decision time, response time, and success rate for M_2 and M_3 over the other models.

A Kruskal-Wallis test indicated a significant effect across the four models for success rate (*SR*), $\chi^2(3, N = 172800) = 7601.2, p < 0.001$, decision time (*DT*), $\chi^2(3, N = 172800) = 15532, p < 0.001$, and response time (*RT*), $\chi^2(3, N = 80528) = 8642.9, p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed a significant effect for success rate ($p < 0.001$) when comparing the first two models, M_0 and M_1, to the interaction delay models, M_2 and M_3. The same test revealed a significant effect between all models for decision time ($p < 0.001$) and response time ($p < 0.001$).

A detailed comparison of the influence of optimal site cost and the time of the site value change (t^F) on each model's performance is provided in Fig. III.10. Fig. III.10a demonstrates the influence of both t^F and optimal site cost on the success rates of the models. M_1 and M_2 are omitted, as no significant effects were observed between those models and M_0 and M_3. M_0 achieved the highest accuracy when responding to changes during the low cost scenario for all values of t^F. M_3 achieved similar success rates to M_0 for the same scenario for $t_F <= 50$s, but was less likely to choose the optimal site, as t^F increased. The models' success rates were similar for the moderate scenario and decayed as t^F increased. The most notable difference between the models was their performance in the high cost scenario. M_0 achieved a success rate of approximately 0.6 when no value change occurred ($t^F = 0$) and its success rate dropped rapidly as t^F increased. M_3's success rates in the high and moderate cost scenarios were nearly equal.

Decision times for different optimal site costs are shown in Fig. III.10b. M_0's and M_1's decision times increased as optimal site cost increased. Lower decisions times were achieved for the easier decisions (low cost) and higher decision times for more difficult, high cost decisions. M_2 was the slowest model overall. M_3 was slower in the low cost scenario, but achieved comparable decisions times to those of M_0 in the moderate and high cost scenarios.

The mean response times for all models in the low cost and high cost scenarios as t^F increases are displayed in Figs. III.10c and III.10d, respectively. M_0 and M_1 responded rapidly when the closer site became the optimal site as a result of a value change at time t^F. Response times for M_0 and M_1 in the high cost scenario, as shown in Fig. III.10d, were much higher than in the low cost scenario. M_2 and M_3 were slower to respond to changes in site value in the low cost scenario, but responded consistently regardless of the optimal site's distance to the collective deliberation area.

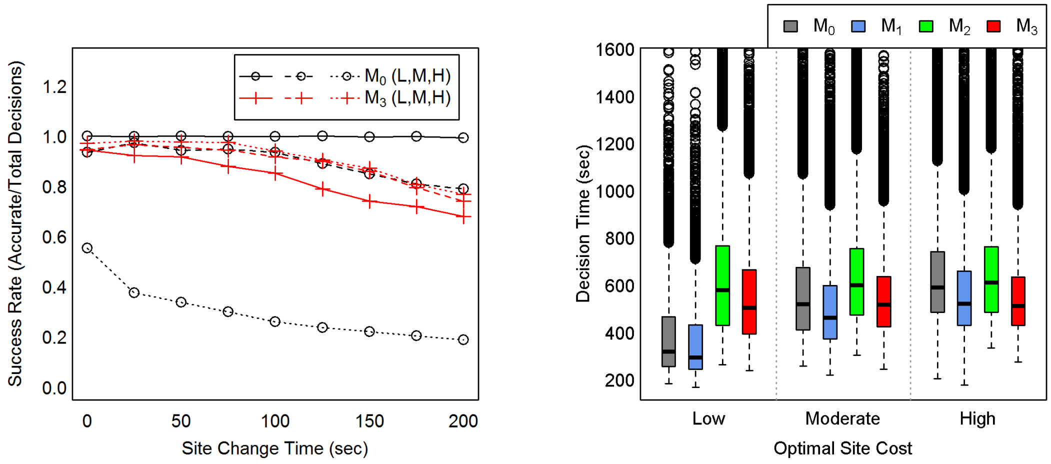

(a) Success Rate for Low-L, Moderate-M and High-H Cost Scenarios

(b) Decision Time for Low-L, Moderate-M and High-H Cost Scenarios

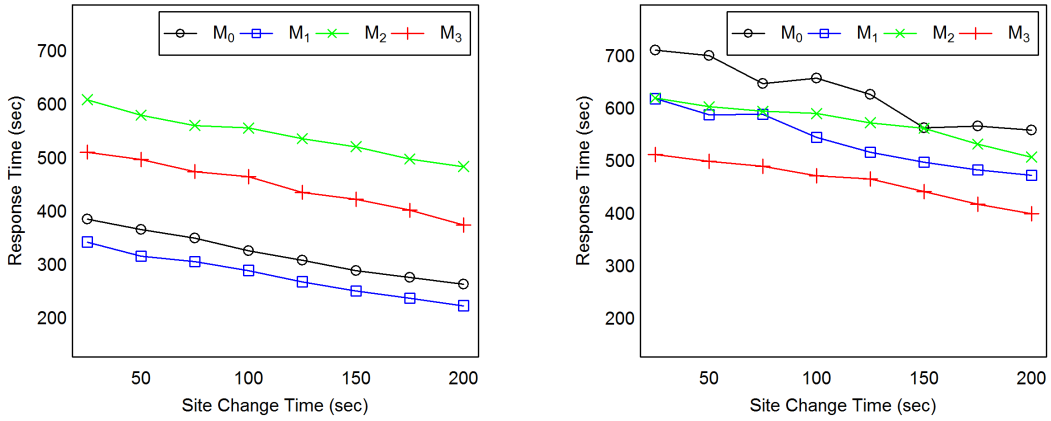

(c) Response Time for Low Cost Scenario: $d_{opt} < d_{med}$

(d) Response Time for High Cost Scenario: $d_{opt} > d_{med}$

Figure III.10: A detailed comparison of the models according to success rate (a), decision time (b), and response time (c and d). Note: Fig. III.10b omits the slowest 1026 trials (0.6% of total trials) with Decision Times exceeding 1600s (M_0-260, M_1-140, M_2-554, M_3-72) in order to facilitate viewing the primary results.

III.3.2.3 Discussion

The results provide further support for hypothesis H_1 (see Chapter III.3.1) by showing that combining frequency modulation and interaction delays improves site selection accuracy in decision problems influenced by negative environmental bias. While, the evaluated models were less accurate, in general, when site values changed later in the decision making process, they demonstrated different behavior patterns in response to dynamic site value changes. The original model's accuracy and response time was strongly influenced by environmental bias and depended upon the locations of the considered sites. Nearby sites had higher discovery rates, because they were easier to locate. Robots supporting nearby sites had shorter latency periods, which increased the frequency of their recruitment and inhibition attempts. Optimal sites in high cost scenarios were harder to find and had longer latency periods than closer sites, which resulted in lower effective recruitment rates and lower success rates.

Hypothesis H_3, which predicted uniform responsiveness to site value changes by the bias reducing models was supported. The model combining both extensions maintained consistent response times, decision times, and success rates that were independent of the evaluated optimal site cost. This consistency resulted in modest decision time improvements and significant success rate improvements during the high optimal site cost scenarios that challenged the original model. Interaction frequency modulation only reduced decision times modestly and had a negligible effect on decision accuracy. While improved decision times are beneficial and enabled the model with both extensions to achieve decision times close to the original model, increasing interaction frequency with site cost was expected to improve success rate more than was observed.

The models with interaction delays (M_2 and M_3) experienced slightly reduced performance in the low optimal site cost scenario, which exposes an important trade-off associated with this extension. The semi-latent control states that enable interaction delays have exit probabilities dependent on a designer-specified maximum search range. Interaction delays were effective in permitting the collective to choose high cost optimal sites, but the delays also slightly slowed the collective's response to changes in the values of the closer sites. Increasing the collective's maximum range is expected to further reduce the collective's responsiveness to value changes for sites close to the deliberation area.

Consistent performance, despite the existence of a biasing factor, indicates the model's ability to ignore aspects of the environment when choosing a site. Collectives that are incorporated into larger efforts, particularly under the supervision of human operators, will be influenced by a variety of environmental biases that will alter the collective's behavior based on its location and the relative locations of the sites it considers. When the collective must make decisions despite a particular environmental bias, the collective's inability to ignore the bias will reduce its performance and increase the challenge of supervising entities, including humans, to predict and understand the collective's expected behavior.

The dynamic site values experiment demonstrates the difficulty of answering research question R_3, since collective responsiveness to site changes depends upon the values and locations of the sites considered, the differences between the considered site values and the timing of a change in site value during the decision making process. The original model, M_0, responded to site changes differently depending upon the location of the sites with respect to the decision making hub. The original model strongly favored closer sites and responded quickly for scenarios in which the closer site changed value, but the original model was much less responsive to value changes with sites located further from the decision making hub. The models that incorporated interaction delays (M_2 and M_3) were consistent in their response, regardless of the location of the sites. Consistent performance in the collective best-of-*n* decision process is important for predicting collective decision times and accuracy in a variety of scenarios. Future human-collective teams that attempt to solve best-of-*n* problems will need to share decision making tasks when neither the human, nor the collective has global knowledge of the environment (Crandall et al., 2017). The consistency exhibited by the interaction delay models (M_2 and M_3) is expected to improve the ability for the human operator to predict the collective's future state and reduce the need for the human to intervene in the event of collective decision errors.

Both the static and dynamic site value experiments evaluated implicit collective action selection scenarios, but did not evaluate each model's ability to detect when a decision has been reached and respond to that decision. Only explicit collective action selection enables a collective to recommended future behavior decisions to the human-collective team, without simultaneously implementing these decisions. The remaining two robotic collective experiments address research question R_4, by examining these models with quorum detection extensions that enabled explicit collective action selection in single (see Chapter III.3.3) and multiple sequential (see Chapter III.3.4) decision scenarios.

III.3.3 Experimental Analysis of Quorum Detection

The quorum detection mechanism described in Chapter III.2.3 is only effective if it enables individual agents to estimate the presence of a quorum with reasonable accuracy. The purpose of the quorum detection experiment is to address research question R_4 by evaluating the collective action selection process's (Chapter III.2.3) ability to respond to site selection by halting the deliberation process. The majority and *k*-unanimity quorum detection mechanisms are combined with the persistent and episodic queuing techniques described in Chapter III.2.3 and are evaluated against the original model (Reina et al., 2015b). Two hypotheses are evaluated. The first hypothesis, $H_{4.1}$, is that the episodic models will be faster and less accurate than the persistent models. The second hypothesis, $H_{4.2}$, is that the unanimous threshold models, called the *k*-unanimity models, will be slower and more accurate than the models using majority quorum thresholds.

III.3.3.1 Experimental Design

Five models, which are summarized in Table III.4, were evaluated in a series of equidistant site experiments. The first model, M_0, was Reina et al.'s (2015b) original model used in the static and dynamic site value experiments. Model M_4 extended M_0 with a persistent quorum detection mechanism and majority quorum rule with a threshold $T = 0.75$ (see Chapter III.2.3). Model M_{4E} was M_4 with an episodic queue that was cleared each time agents changed state. Model M_5 used a persistent queue with a strict, *k*-unanimity, threshold ($T = 1.0$). Finally, Model M_{5E} was M_5 with an episodic queue. The arena, shown in Fig. III.11, was similar to the prior experiments, (see Chapter III.3.1 and III.3.2). Agent movement, sensory, and communication

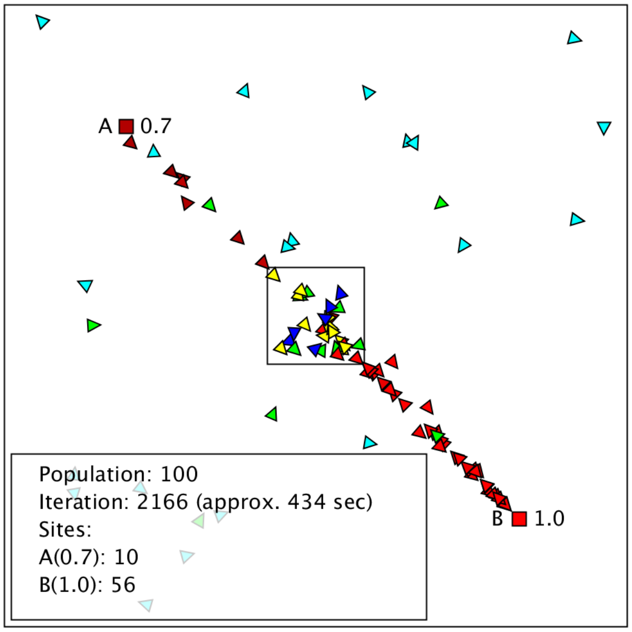

Population: 100
Iteration: 2166 (approx. 434 sec)
Sites:
A(0.7): 10
B(1.0): 56

Figure III.11: A screenshot of a trial in which the collective must choose between an optimal and mediocre site. The collective has just detected a quorum for site *B*, indicated by the dark blue agents in the central decision making hub.

capabilities were identical to previous experiments. The secondary independent variables included site value, in the range $[0.7, 1.0]$, incremented by 0.1, and collective population, in the range $[100, 500]$, incremented by 100. Each model was evaluated for 50 trials of each combination of secondary independent variables in two site decision problems, where the sites were equidistant, placed at 2m from the nest area, as shown in Fig. III.11. The evaluation metrics were success rate, decision time, quorum accuracy, and split decision rate. Success rate was the ratio of accurate decisions to total decisions for each problem scenario (Reina et al., 2015b). Decision time was the time required for the collective to sense a 75% majority opinion (Reina et al., 2015b). Quorum accuracy was the ratio of the population percentage supporting a site at the time a quorum was sensed to the desired 75% threshold. Finally, the split decision rate was the ratio of split decisions to total decisions for each problem scenario. M_0 was evaluated using success rate and decision time, which were

Table III.4: Evaluated Models

Model	Description
M_0	Original model
M_4	Majority Quorum with Persistent Queue
M_{4E}	Majority Quorum with Episodic Queue
M_5	*k*-Unanimity Quorum with Persistent Queue
M_{5E}	*k*-Unanimity Quorum with Episodic Queue

externally observed, but the remaining metrics were not applicable to the model, as it did not sense collective decisions.

III.3.3.2 Results

The descriptive statistics for success rate, decision time, and quorum accuracy are presented in Table III.5. The 0.3% of trials ending in a split decision are omitted from the table, but are presented in Fig. III.12(d). All evaluated models achieved similar mean, median, standard deviation (SD), minimum (Min), and maximum (Max) success rate values. The persistent models (M_4 and M_5) resulted in longer decision times and higher quorum accuracies than the episodic models (M_{4E} and M_{5E}). M_4 achieved the highest quorum accuracy, while M_{5E} made the fastest decisions with the lowest quorum accuracy.

A Kruskal-Wallis test did not find significant effects for success rate. Significant effects were observed for decision time, $\chi^2(4, N = 20000) = 2210.5$, $p < 0.001$ and quorum accuracy, $\chi^2(4, N = 20000) = 2789$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for decision time ($p < 0.01$) and quorum accuracy ($p < 0.001$) between all models.

A detailed comparison of model performance with respect to collective population size is presented for success rate (Fig. III.12(a)), decision time (Fig. III.12(b)), quorum accuracy (Fig. III.12(c)), and split decision rate (Fig. III.12(d)). Success rates for all models improved as collective population increased, although success rates leveled off when the collective population reached 400 agents. The evaluated queue length is the expected cause of this observation.

The persistent quorum sensing models were significantly slower to sense a decision than the episodic models, as shown in Fig. III.12(b), although all models improved with an increase in population size.

Table III.5: Descriptive Statistics: Success Rate (*SR*), Decision Time (*DT*), and Quorum Accuracy (*QA*).

Model	Statistic	SR (%)	DT (minutes)	QA
M_0	Mean (SD)	0.898 (0.10)	16.01 (10.1)	N/A
	Median	0.92	13.28	
	Min (Max)	0.62 (1.0)	6.5 (134.9)	
M_4	Mean (SD)	0.881 (0.11)	15.3(3.3)	0.99 (0.19)
	Median	0.92	14.4	1.04
	Min (Max)	0.56 (1.0)	9.7 (37.4)	0.17 (1.3)
M_{4E}	Mean (SD)	0.876 (0.10)	9.3 (4.2)	0.56 (0.11)
	Median	0.90	8.13	0.56
	Min (Max)	0.64 (1.0)	3.1 (50.7)	0.13 (0.93)
M_5	Mean (SD)	0.889 (0.10)	19.2(6.6)	1.2(0.1)
	Median	0.92	18.0	1.15
	Min (Max)	0.64 (1.0)	10.78 (92.81)	0.56 (1.29)
M_{5E}	Mean (SD)	0.890 (0.11)	14.4(6.9)	0.91(0.10)
	Median	0.91	12.5	0.92
	Min (Max)	0.6 (1.0)	4.68 (69.3)	0.44 (1.17)

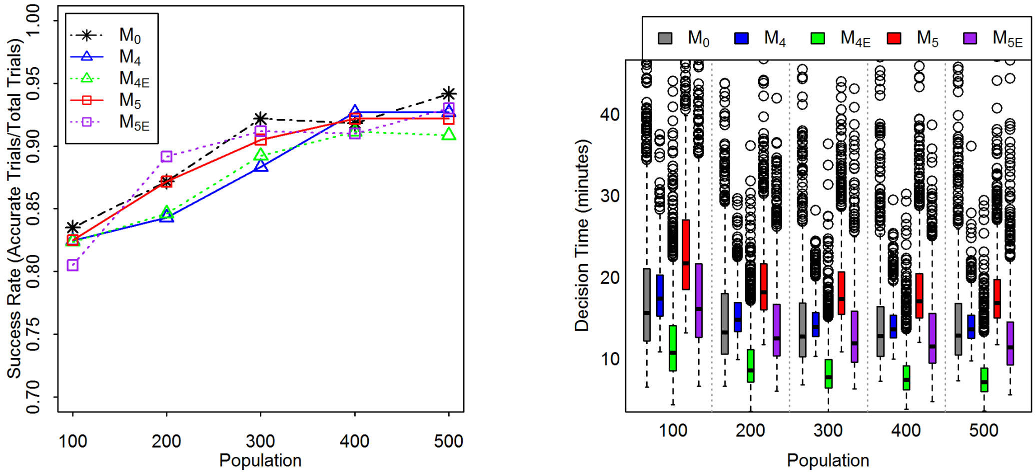

(a) Success Rate

(b) Decision Time

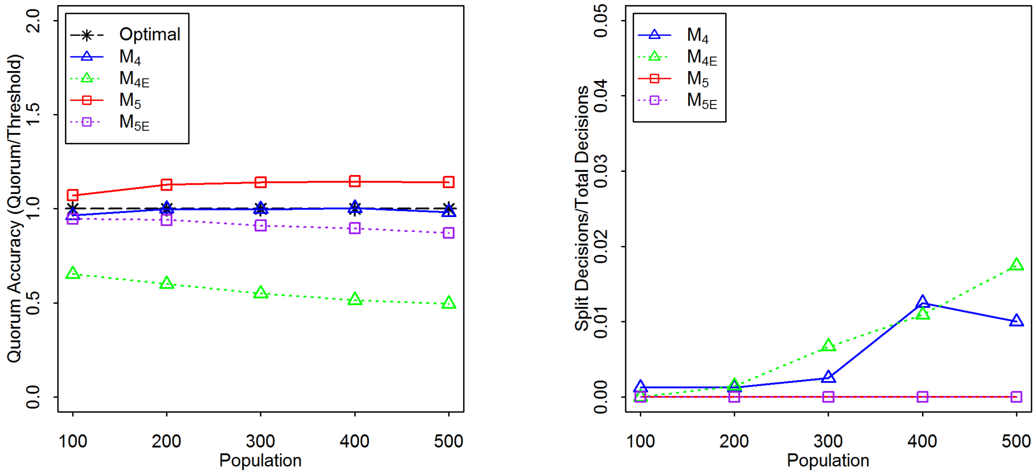

(c) Quorum Accuracy

(d) Split Decision Rate

Figure III.12: The performance of the five evaluated models are compared based on success rate a and decision time (b) across the evaluated population sizes. All models except M_0 are also compared by quorum accuracy (c) and (d) across the evaluated population sizes.

Episodic queues resulted in decreases of 40% and 25% in decision time for the majority rule and k-unanimity rule, respectively, without a reduction in success rates.

M_{4E}'s low decision time was due to its tendency to sense a quorum far before the other models, as shown in Fig. III.12(c). The remaining models sensed a quorum very near to when the majority opinion reached the 75% threshold, and were not influenced by increasing the collective population size. M_{5E} exceeded expectations by achieving fast decision times with corresponding high quorum accuracy.

Split decision rates proved to be an important discriminator between the quorum sensing models. While Figs. III.12(a-c) ignore the 0.3% of overall decisions that resulted in split decisions, Fig. III.12(d) demonstrates that all split decisions occurred with the majority rule models, M_4 and M_{4E}. Split decisions increased with population size for both models, indicating that there is an increasing likelihood of agents sensing a false quorum after receiving most of their messages from a minority of the population. Models M_5 and M_{5E} did not experience a single split decision during the experiment.

The highest performing quorum sensing models were M_4 and M_{5E}. Both models made and detected site selections with near optimal quorum accuracy and decision times comparable to the original model (M_0). M_{4E} experienced a 43% drop in quorum accuracy compared to M_4, while M_{5E} was just 23% lower than M_5 for the same frequency. Further, M_4 enabled detection of the desired threshold ($T = 0.75$), while experiencing a low percentage of split decisions. M_{5E} avoided split decisions and achieved similar decision times to M_4.

III.3.3.3 Discussion

The results demonstrate that both the majority and *k*-unanimity quorum detection rules enable the collective to achieve value-sensitive consensus decisions in collective site selection problems. The results also indicate that Parker and Zhang's (2009) quorum sensing mechanism is applicable to a large collective restricted to communication in a decision making hub. Finally, the performance of the persistent majority rule and episodic *k*-unanimity rule quorum mechanisms demonstrate that collective decision detection does not necessarily increase the decision time of Reina *et al.*'s underlying model.

The first hypothesis, $H_{4.1}$, which stated that persistent queues were expected to be more accurate and slower than episodic queues was only partially supported. Both episodic models were faster than the persistent models, but each was significantly more accurate than expected. While reduced decision speed is often desirable, collective commitment to an opinion ahead of achieving a true majority is expected to be detrimental in scenarios with dynamically changing site values. A key attribute of the underlying decision strategy is its value-sensitivity, which is the collective's ability to remain in a deliberative state in the presence of poor quality sites, until acceptable sites are discovered (Pais et al., 2013). Early quorum detection increases the probability that the collective commits to a poor site, which reduces the system's value-sensitivity.

The results also partially supported hypothesis $H_{4.2}$ which stated the majority rule models were expected to be faster and less accurate than the *k*-unanimity rule models. The majority rule models achieved shorter decision times with similar decision accuracy when compared to the *k*-unanimity rule models, but they were also more likely to result in a split decision. Split decisions were primarily a problem due to the current absorbing commitment states. An alternative solution is to permit competition between the commitment states, achieved through application of a consensus decision strategy, such as the *k*-unanimity rule (Scheidler et al., 2016). Examination of this secondary consensus process will establish if the decision time gained by a less stringent quorum detection rule is truly advantageous.

The quorum detection experiment begins to address research question R_4 by evaluating the collective's ability to make a best-of-*n* decision and transition its behavior based on quorum detection (e.g., stop deliberating between available sites). This experiment indicates that the presented solution enables the collective to simultaneously solve the best-of-*n* problem and sense a quorum without reducing the accuracy of the initial decision in the evaluated cases. A comparison between the decision times of the lower quorum threshold models (M_4 and M_{4E}) to the higher threshold models (M_5 and M_{5E}) demonstrates clearly that increasing quorum threshold values increased decision time and decreased the probability that the collective splits for the evaluated population sizes. The former result is not surprising (Parker and Zhang, 2009), but simultaneous quorum detection is a potentially catastrophic problem in scenarios that require the collective to move the location of the decision making hub after a decision. Split decisions under these conditions result in the collective fragmenting as portions of it attempt to set up different decision making hubs in new locations. The difference in performance between the models using persistent (M_4 and M_5) and episodic (M_{4E} and M_{5E}) queues demonstrated that rapid quorum detection did not necessarily reduce decision accuracy, but did influence quorum accuracy.

Addressing the remainder of R_4 requires that the collective be evaluated as it completes multiple decisions in a row. Sequential decisions require several different behavior changes, as the collective commits to an action, initiates the action, and completes the action. Split decisions during sequential decision making are likely to occur when the collective simultaneously detects quorums for multiple actions, as demonstrated in this experiment, or when the collective transitions to a new behavior and leaves individual robots behind in a previous state within the decision making process.

III.3.4 Experimental Analysis of Sequential Action Selection

The explicit collective action selection model (Chapter III.2.3) requires the collective to alternate action decisions and executions in a sequence of problems. The quorum sensing evaluation indicated that increasing the frequency of agent voting messages decreased quorum accuracy, but that effective quorum detection was possible at the higher frequency with an increased quorum threshold. The purpose of the sequential decisions experiment was to address research question R_4 by evaluating the presented explicit collective action selection approach in scenarios that required the completion of a sequence of action selection problems. Four different models were evaluated and distinguished by the use of bias reduction mechanisms and the use of either a persistent or episodic voting queue. Additionally, the models were tested in four different quorum noise scenarios in order to evaluate the effectiveness of different noise reduction techniques. The first hypothesis, $H_{4.3}$, was that the persistent queue techniques enable more sequential decision making than the episodic models. The second hypothesis, $H_{4.4}$, was that quorum noise reduction techniques improve the speed and completion rate of all evaluated models.

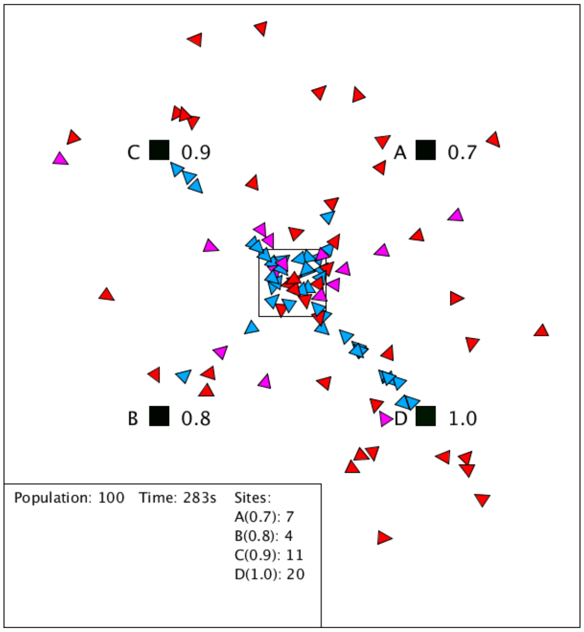

Figure III.13: An example of the beginning of the first of four decisions the collective made during a sequential trial. The collective chose each site in a sequence determined by its collective best-of-*n* model.

III.3.4.1 Experimental Design

The arena dimensions and agent capabilities for the sequential action selection experiment were identical to the previously presented experiments; however, the collective search space contained four sites. Unlike prior experiments, a single arrangement of sites and values, shown in Figure III.13, was used to evaluate the collective's ability to complete a sequence of action selections. The sites were located on four corners of a square, with each site located 1.25m from the initial, central location of the collective's decision making hub. The sites had static values in the range $[0.7, 1.0]$, arranged as shown in Figure III.13. The collective was required to choose each site in a sequence determined by its collective best-of-*n* model. The first decision required the collective to choose the best of four sites, detect its decision, and move its decision making hub to the chosen site. The second decision required the collective to choose between the remaining three sites, which was repeated for the last two decisions. Due to the arrangement of sites in a square, the distances between the sites and the collective decision making hub were always within the set $\{1.25, 1.8, 2.5\}$m.

Table III.6: Evaluated Models

Model	Description
M_4	Action Selection with Persistent Queue
M_{4E}	Action Selection with Episodic Queue
M_6	M_4 with Bias Reduction
M_{6E}	M_{4E} with Bias Reduction

The primary independent variable was the evaluated model. The four evaluated models are listed in Table III.6. Models M_4 and M_{4E} were identical to those used in the quorum detection experiment (see Chapter III.3.3) and used persistent and episodic queues, respectively. Model M_6 and M_{6E} added interaction delay and interaction frequency modulation to M_4 and M_{4E}, respectively. The secondary independent variables were the collective population size, $N \in \{100, 300, 500\}$, quorum thresholds $T \in \{0.8, 1.0\}$, and four different quorum noise scenarios. Each primary independent variable was paired with each combination of secondary independent variables and evaluated via 50 trials. Quorum detection is susceptible to noise that results from duplicate messages being received from neighboring agents, as described in Chapter III.2.3. The baseline, unfiltered quorum noise scenario (UQ) made no attempt to filter quorum messages and individuals were susceptible to receiving duplicate messages. The delayed interaction scenario (DI) reduced agent interaction probabilities by half in order to evaluate Parker and Zhang's (2010) recommendation that increasing delays between interactions passively reduces the likelihood of duplicate messages. The weak quorum filtering scenario (WQ) required agents to generate a message identifier in the range $[0, k-1]$ each time they became interactive and include the identifier in voting messages. Agents that received a message with an identifier that was already in its queue dropped the new message. Finally, the strong quorum message filtering scenario (SQ) required agents to generate a random 32-bit number each interaction period and ignore duplicate messages.

The metrics were success rate, action completion time, quorum accuracy, and trial completion rate. Success rate, as in prior experiments, was the ratio of accurate decisions to the total decisions made by the collective. Action completion time was the time required for the collective to choose a site, move to the site, and achieve a quorum at the new location. Quorum accuracy was the ratio of the actual number of agents in a particular quorum sensing state at the time quorum was detected to the optimal number of agents determined by the product of the quorum threshold T and the agent population N. Finally, trial completion rate was the ratio of completed trials to the total trials. Collectives that failed to complete the required four decisions were considered to have failed the trial. Failed trials occurred when the collective was either unable to make its last decision within six simulated hours, or if the collective split during the trial. Split collectives occurred if the collective either committed simultaneously to two different sites or if the collective left agents behind at a prior nest location.

III.3.4.2 Results: Unfiltered Quorum Noise Scenario (UQ)

Each model was compared according to the descriptive statistics for success rate (SR), action completion time (AT), quorum accuracy (QA), and trial completion rate (TC). The descriptive statistics for the unfiltered queue scenario in which agents do not filter quorum messages and do not delay between interactions, are provided in Table III.7. M_6^{UQ} and M_{6E}^{UQ} achieved the highest success rates, which indicates that when these

Table III.7: Unfiltered Queue (UQ): Descriptive Statistics by Model and Metric.

Model	Statistic	SR	AT	QA	TC
M_4^{UQ}	Mean (SD)	0.64 (0.48)	11.79 (3.81)	0.82 (0.34)	0.8 (0.4)
	Median	1	11.54	0.91	1
	Min (Max)	0 (1)	5.5 (54.26)	0.01 (1.24)	0 (1)
M_{4E}^{UQ}	Mean (SD)	0.6 (0.49)	3.15 (1.77)	0.19 (0.11)	0 (0)
	Median	1	2.52	0.16	0
	Min (Max)	0 (1)	0 (12.18)	0 (0.58)	0 (0)
M_6^{UQ}	Mean (SD)	0.92 (0.27)	20.78 (6.85)	0.92 (0.14)	0.63 (0.48)
	Median	1	19.06	0.93	1
	Min (Max)	0 (1)	0 (120.4)	0.36 (1.18)	0 (1)
M_{6E}^{UQ}	Mean (SD)	0.83 (0.38)	8.86 (4.72)	0.28 (0.08)	0.02 (0.13)
	Median	1	7.72	0.29	0
	Min (Max)	0 (1)	0 (47.37)	0 (0.52)	0 (1)

models made decisions, the decisions were likely to be accurate. The action completion times achieved by M_{4E}^{UQ} and M_{6E}^{UQ} were very low, indicating that each of these models made fast decisions, but also fragmented during the trial. M_{4E}^{UQ} did not complete any trials and M_{6E}^{UQ} completed only a negligible number of trials. M_6^{UQ} was the slowest and most accurate model, but M_4^{UQ} achieved the highest trial completion rate of all the models.

A Kruskal-Wallis test indicated significant effects for success rate, $\chi^2(3, N = 13155) = 1027$, $p < 0.001$, action completion time, $\chi^2(3, N = 13155) = 8298.1$, $p < 0.001$, quorum accuracy, $\chi^2(3, N = 13155) = 8893$, $p < 0.001$ and trial completion $\chi^2(3, N = 4500) = 2245.1$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for action completion time ($p < 0.001$) and quorum accuracy ($p < 0.001$) between all models. No significant effects were observed for success rate between models M_4^{UQ} and M_{4E}^{UQ}, but significant effects were observed in pairwise comparisons between the rest of the models ($p < 0.001$). Finally, no significant effects were observed for trial completion between the episodic models, M_{4E}^{UQ} and M_{6E}^{UQ}, although the remaining pairwise comparisons between the models revealed significant effects for trial completion ($p < 0.001$).

The models in scenario UQ are compared in Figure III.14 for success rate (Fig. III.14a), action completion time (Fig. III.14b), quorum accuracy (Fig. III.14c), and trial completion (Fig. III.14d). Models M_6^{UQ} and M_{6E}^{UQ} had higher success rates for each decision when compared to models M_4^{UQ} and M_{4E}^{UQ}, indicating each model's resilience to site distance from the decision making hub. Model M_4^{UQ} consistently faced a difficult decision after choosing its first site and was usually unable to choose the higher valued site that was further from its decision making hub than the other sites. Model M_{4E}^{UQ}'s apparent accuracy in the second decision resulted from making poor first decisions and facing an easier second decision. M_{4E}^{UQ} failed to make a fourth decision for every single trial, which is indicated by the model's missing data point. The

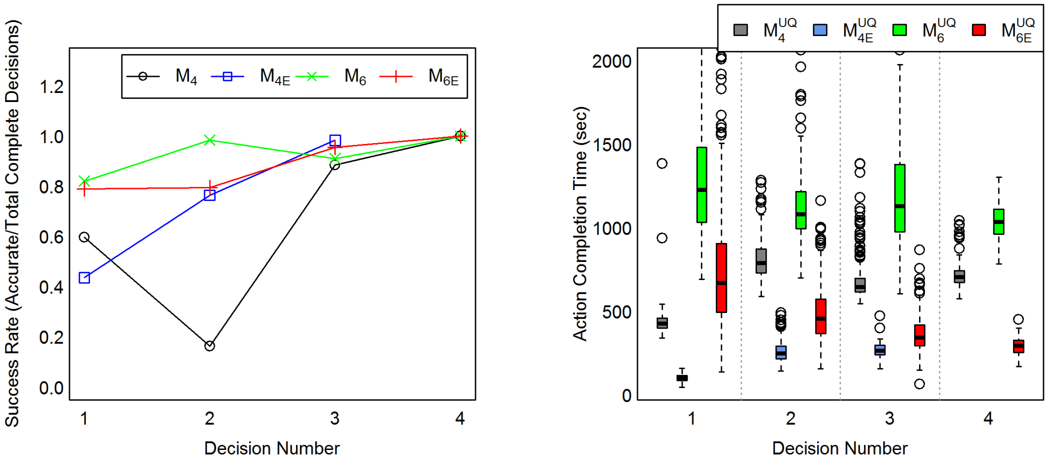

(a) Success Rate

(b) Action Completion Time

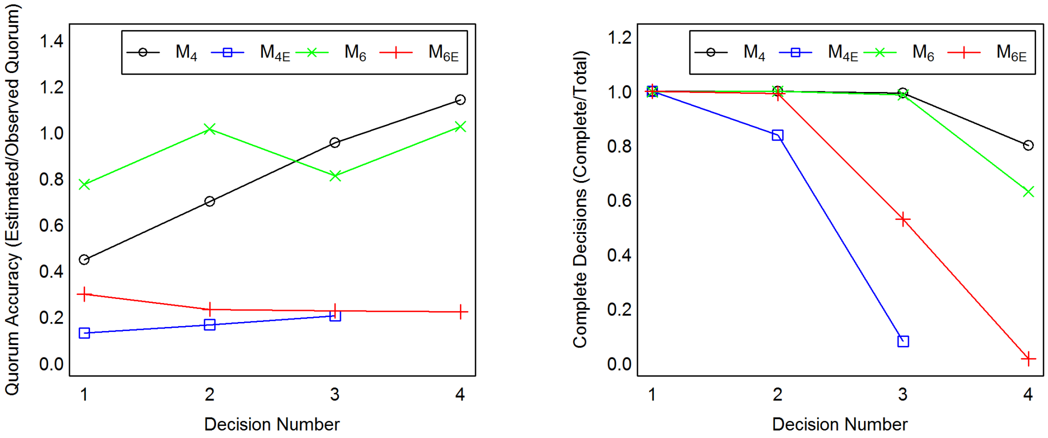

(c) Quorum Accuracy

(d) Trial Completion Rate

Figure III.14: UQ (Unfiltered Queue): The models are compared by success rate (Fig. a), action completion time (Fig. b), quorum accuracy (Fig. c), and trial completion rate (Fig. d). Note: Fig. (b) cuts off the slowest 254 trials (3% of the total trials) with Action Completion Times exceeding 2000s (M_4-21,M_{4E}-0,M_6-218,M_{6E}-15) in order to facilitate viewing the primary results.

action completion times differed significantly between the models, as shown in Fig. III.14b, with the models performing differently between the first decision and the remaining decisions. M_4^{UQ} and M_{4E}^{UQ} made rapid first decisions, but made slower decisions for the remainder of the trial, with M_{4E}^{UQ} clearly making the fastest decisions over all. The interaction delay models were each slower with the first decision, but made faster decisions for the remaining decisions in each trial. M_{6E}^{UQ} achieved faster action completion times, when it made decisions, than M_4^{UQ}. The persistent queue models, M_4^{UQ} and M_6^{UQ} achieved much higher quorum accuracy than the remaining models in scenario UQ. The quorum accuracy of M_4^{UQ} tended to increase as it progressed through the trials, while M_6^{UQ} was noticeably more stable. Finally, the trial completion rates for each model are presented in Fig. III.14d. The episodic queue models (M_{4E}^{UQ} and M_{6E}^{UQ}) successfully completed a majority of the first two decisions, but were unable to complete the remainder of the site decisions for most trials. The persistent queue models consistently made three decisions, with M_4^{UQ} completing significantly more trials than M_6^{UQ}.

III.3.4.3 Results: Delayed Interaction Scenario (*DI*)

The Delayed Interaction scenario (*DI*) descriptive statistics are presented in Table III.8. Longer interaction times resulted in increased success rates, action completion times, and quorum accuracies across all four models. The trial completion rates for the persistent queue models (M_4^{DI} and M_6^{DI}) were slightly higher than the completion rates for M_4^{UQ} and M_6^{UQ} in the unfiltered queue scenario (UQ). The models otherwise performed similarly to the UQ scenario. M_6^{DI} and M_{6E}^{DI} achieved the highest accuracy. The episodic queue models (M_{4E}^{DI} and M_{6E}^{DI}) made faster decisions than the persistent queue models (M_4^{DI} and M_6^{DI}), but also failed to complete trials. The results indicate that increasing the delay period alone improved both quorum accuracy and trial completion, which supports the unevaluated recommendations from prior work (Parker and Zhang, 2010). The small action completion time increase observed between M_6^{UQ} and M_6^{DI} indicate that the model imposed delays during the UQ scenario that approached those used during the *DI* scenario.

A Kruskal-Wallis test indicated significant effects for success rate, $\chi^2(3, N = 11718) = 963.34$, $p < 0.001$, action completion time, $\chi^2(3, N = 11718) = 6684.1$, $p < 0.001$, quorum accuracy, $\chi^2(3, N = 11718) = 8178.8$, $p < 0.001$ and trial completion $\chi^2(3, N = 2037) = 2037$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for action completion time ($p < 0.001$), quorum accuracy ($p < 0.001$), and trial completion ($p < 0.001$) between all models. Significant effects for success rate were not observed between the interaction delay models (M_6^{DI} and M_{6E}^{DI}), but significant effects were observed in the remaining pairwise comparisons ($p < 0.001$).

Table III.8: Delayed Interaction (*DI*): Descriptive Statistics by Model and Metric.

Model	Statistic	SR	AT (min)	QA	TC
M_4^{DI}	Mean (SD)	0.66 (0.47)	14.68 (4.33)	1.04 (0.2)	0.9 (0.3)
	Median	1	13.65	1.13	1
	Min (Max)	0 (1)	6.98 (59.74)	0.12 (1.24)	0 (1)
M_{4E}^{DI}	Mean (SD)	0.64 (0.48)	4.8 (1.96)	0.31 (0.12)	0 (0)
	Median	1	4.5	0.3	0
	Min (Max)	0 (1)	0 (15.07)	0 (0.64)	0 (0)
M_6^{DI}	Mean (SD)	0.92 (0.27)	20.97 (6.2)	0.93 (0.15)	0.71 (0.45)
	Median	1	19.53	0.94	1
	Min (Max)	0 (1)	0 (66.04)	0.44 (1.24)	0 (1)
M_{6E}^{DI}	Mean (SD)	0.87 (0.34)	11.06 (5.53)	0.37 (0.09)	0.01 (0.1)
	Median	1	9.69	0.38	0
	Min (Max)	0 (1)	0 (44.57)	0 (0.62)	0 (1)

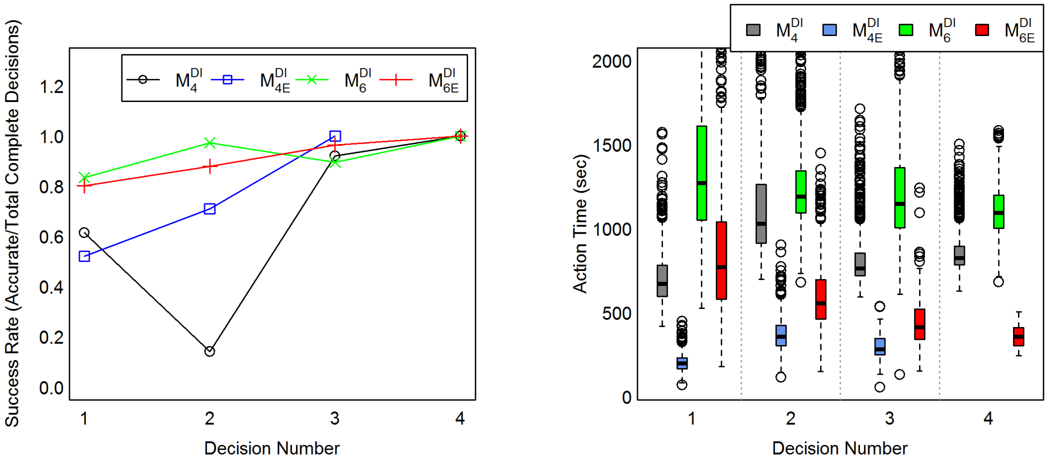

(a) Success Rate

(b) Action Completion Time

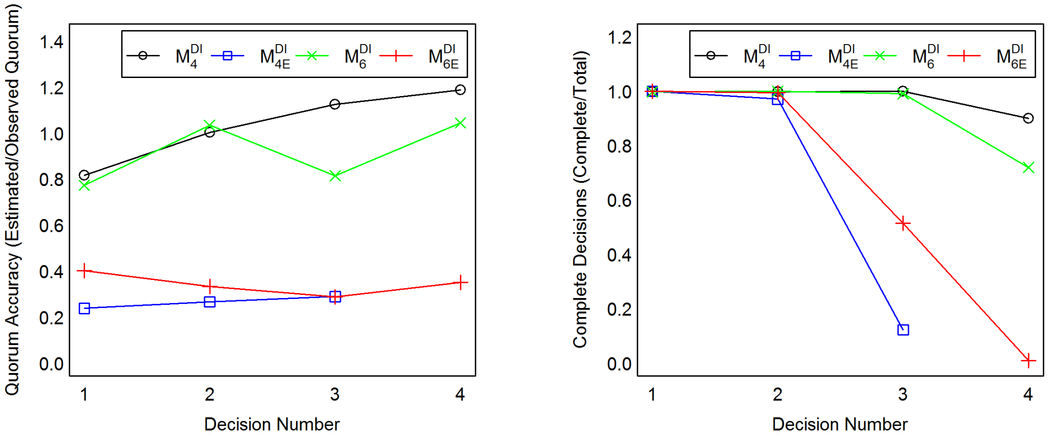

(c) Quorum Accuracy

(d) Trial Completion Rate

Figure III.15: *DI* (Delayed Interaction): The models are compared by success rate (Fig. a), action completion time (Fig. b), quorum accuracy (Fig. c), and trial completion rate (Fig. d). Note: Fig. (b) cuts off the slowest 256 trials (3% of the total trials) with Action Completion Times exceeding 2000s (M_4-35,M_{4E}-0,M_6-292,M_{6E}-29) in order to facilitate viewing the primary results.

The behavior of the four *DI* scenario models is presented in Fig. III.15. Success rates and action completion times increased overall, as shown in Figs. III.15a and III.15b, respectively. Quorum accuracy improved for the M_4^{DI}, M_{4E}^{DI}, and M_{6E}^{DI} models over the UQ models, M_4^{UQ}, M_{4E}^{UQ}, and M_{6E}^{UQ}, respectively (see Figure III.14c). Model M_6^{DI} achieved similar quorum accuracy to M_6^{UQ} (see Figure III.14c). The quorum accuracy improved the most between M_4^{UQ} and M_4^{DI} when compared to the other models. M_4^{DI} maintained a mean 80% quorum accuracy, or greater, for every decision. Doubling the delay period between interactions modestly improved the completion rates for the two persistent queue models (M_4^{DI} and M_6^{DI}), when compared to M_4^{UQ} and M_6^{UQ}, but the adjustment did not noticeably improve the completion rates for the remaining models.

Table III.9: Weak Quorum Filtering (WQ): Descriptive Statistics by Model and Metric.

Model	Statistic	SR	AT	QA	TC
M_4^{WQ}	Mean (SD)	0.67 (0.47)	16.91 (5.54)	1.05 (0.27)	0.97 (0.17)
	Median	1	16.06	1.19	1
	Min (Max)	0 (1)	6.7 (61.02)	0.12 (1.25)	0 (1)
M_{4E}^{WQ}	Mean (SD)	0.63 (0.48)	6.04 (2.41)	0.52 (0.23)	0.88 (0.32)
	Median	1	5.88	0.53	1
	Min (Max)	0 (1)	0 (47.89)	0 (1.04)	0 (1)
M_6^{WQ}	Mean (SD)	0.92 (0.27)	27.89 (7.45)	1.09 (0.08)	0.95 (0.21)
	Median	1	26.16	1.11	1
	Min (Max)	0 (1)	13.82 (107.7)	0.58 (1.23)	0 (1)
M_{6E}^{WQ}	Mean (SD)	0.91 (0.29)	12.51 (5.75)	0.49 (0.12)	0.77 (0.42)
	Median	1	10.98	0.5	1
	Min (Max)	0 (1)	3.32 (63.23)	0 (0.89)	0 (1)

III.3.4.4 Results: Weak Quorum Filtering Scenario (WQ)

The Weak Quorum Filtering (WQ) scenario required individual agents to send interaction identifiers (in the range $[0, k-1]$) with queue messages and to ignore received messages with identifiers that matched those already in the queue. The descriptive statistics are presented in Table III.9. The model success rates were similar to those recorded for the UQ (see Fig. III.14a) and DI (see Fig. III.15a) scenarios, which indicates that even when false quorums are detected as a result of duplicated messages, the collective frequently committed early to the site that it ultimately selected. Action completion times across the models were slower overall, with the persistent queue models (M_4^{WQ} and M_6^{WQ}) experiencing the largest increase in action completion time. Quorum accuracy and trial completion also increased for all models. This scenario was the first in which the episodic queue models (M_{4E}^{WQ} and M_{6E}^{WQ}) completed a majority of their trials.

A Kruskal-Wallis test indicated significant effects for success rate, $\chi^2(3, N = 17158) = 1637.4$, $p < 0.001$, action completion time, $\chi^2(3, N = 17158) = 11574$, $p < 0.001$, quorum accuracy, $\chi^2(3, N = 17158) = 11601$, $p < 0.001$ and trial completion $\chi^2(3, N = 4423) = 316.94$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for action completion time ($p < 0.001$) between all models. Significant effects for success rate were not observed in pairwise comparisons between M_6^{WQ} and M_{6E}^{WQ}, or between M_4^{WQ} and M_{4E}^{WQ}, although significant effects were observed in the remaining model comparisons ($p < 0.001$). The analysis revealed significant effects for quorum accuracy between M_4^{WQ} and M_6^{WQ} ($p < 0.01$) and in all other model comparisons ($p < 0.001$). Finally, significant effects were observed for trial completion between M_4^{WQ} and M_{4E}^{WQ} ($p < 0.05$) and between M_{6E}^{WQ} and all other models ($p < 0.001$).

The addition of weak queue filtering notably improved the ability of the episodic queue models (M_{4E}^{WQ} and M_{6E}^{WQ}) to make sequential decisions, as shown in Fig. III.16. The models achieved success rates (Fig. III.16a) similar to the UQ and DI models with the exception of M_{4E}^{WQ}, which, like M_4^{WQ}, struggled during the second

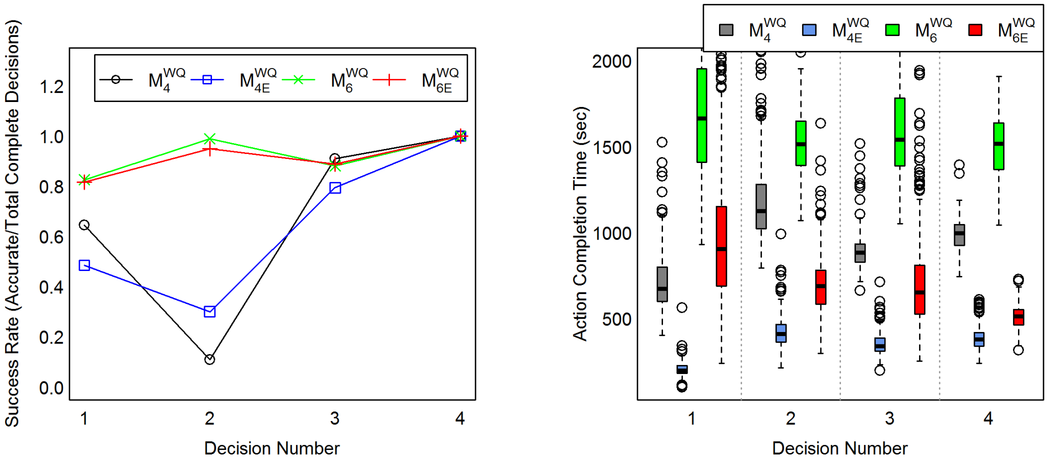

(a) Success Rate

(b) Action Completion Time

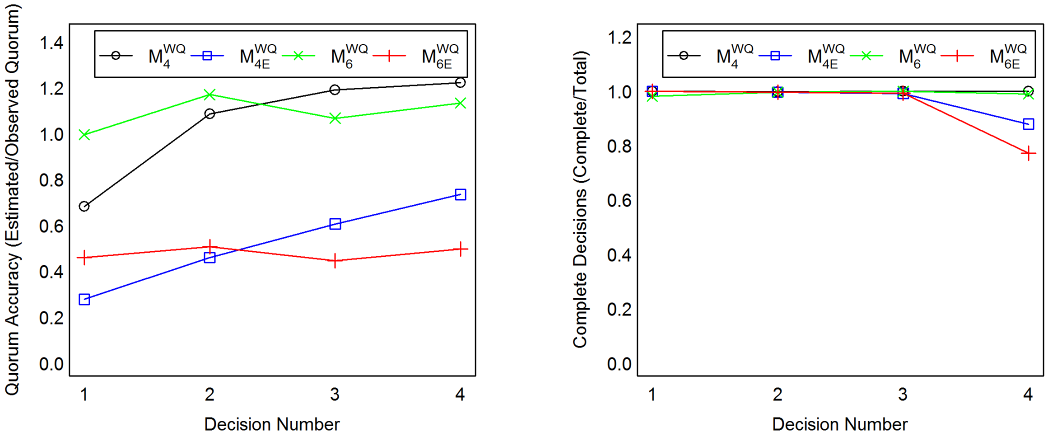

(c) Quorum Accuracy

(d) Trial Completion Rate

Figure III.16: WQ (Weak Filtered Queue): The models are compared by success rate (Fig. a), action completion time (Fig. b), quorum accuracy (Fig. c), and trial completion rate (Fig. d). Note: Fig. (b) cuts off the slowest 1381 trials (15% of the total trials) with Action Completion Times exceeding 2000s (M_4-203,M_{4E}-7,M_6-1113,M_{6E}-58) in order to facilitate viewing the primary results.

decision. Action completion times (Fig. III.16b) were higher for all the WQ models, indicating that the use of weak filtering increased the length of time required to detect a quorum. Quorum accuracy, shown in Fig. III.16c, improved for the episodic queue models (M_{4E}^{WQ} and M_{6E}^{WQ}), as well as for M_6^{WQ}, but M_4^{WQ}'s quorum accuracy for its first decision was slightly lower than the corresponding *DI* scenario model. All four models completed a majority of the trials, as shown in Fig. III.16d, with the episodic queue models completing at least three decisions in almost every trial.

Table III.10: Strong Quorum Filtering (*SQ*): Descriptive Statistics by Model and Metric.

Model	Statistic	SR	AT	QA	TC
M_4^{SQ}	Mean (SD)	0.64 (0.48)	11.74 (3.52)	0.83 (0.35)	0.97 (0.18)
	Median	1	11.56	0.92	1
	Min (Max)	0 (1)	5.54 (38.55)	0.01 (1.24)	0 (1)
M_{4E}^{SQ}	Mean (SD)	0.6 (0.49)	5.3 (2.16)	0.45 (0.19)	0.96 (0.2)
	Median	1	5.02	0.46	1
	Min (Max)	0 (1)	1.48 (20.47)	0.01 (0.95)	0 (1)
M_6^{SQ}	Mean (SD)	0.92 (0.27)	21.29 (6.39)	0.94 (0.17)	0.96 (0.19)
	Median	1	19.53	0.98	1
	Min (Max)	0 (1)	9.86 (71)	0.02 (1.21)	0 (1)
M_{6E}^{SQ}	Mean (SD)	0.91 (0.29)	10.85 (5.17)	0.41 (0.11)	0.95 (0.22)
	Median	1	9.35	0.42	1
	Min (Max)	0 (1)	2.84 (68.35)	0 (0.79)	0 (1)

III.3.4.5 Results: Strong Filtered Quorum Noise Scenario (*SQ*)

The use of strong queue filtering was intended to eliminate false quorum detection that results from duplicative messages. The descriptive statistics for the *SQ* models are provided in Table III.10. The models' success rates were consistent with the *UQ*, *DI*, and *WQ* scenarios. Action completion times for the persistent queue models (M_4^{SQ} and M_6^{SQ}) were similar to those of achieved by M_4^{UQ} and M_6^{UQ} (see Fig. III.14b). Both episodic models made rapid decisions, with M_{4E}^{SQ} achieving the fastest action completion times of the *SQ* models. Quorum accuracies decreased for the episodic models when compared to the weak filtering scenario, although both M_{4E}^{SQ} and M_{6E}^{SQ} completed nearly as many trials as the persistent queue models. Trial completion was reasonably high for all models.

A Kruskal-Wallis test indicated significant effects for success rate, $\chi^2(3, N = 14830) = 1883.8$, $p < 0.001$, action completion time, $\chi^2(3, N = 14830) = 9460.8$, $p < 0.001$, and quorum accuracy, $\chi^2(3, N = 14830) = 6937.5$, $p < 0.001$. No significant effect was observed for trial completion. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for action completion time ($p < 0.001$) and quorum accuracy ($p < 0.001$) between all models. Significant effects were observed for success rate between the interaction delay models (M_6^{SQ} and M_{6E}^{SQ}) and the remaining models ($p < 0.001$).

A detailed comparison of the *SQ* models is provided in Fig. III.17. Success rates for each model by decision were similar to the *WQ* scenario. The *SQ* models' action completion times were faster than the *WQ* models (see Fig. III.16b), with the persistent queue models achieving action completion times similar to the *UQ* scenario (see Fig. III.14b). The use of strong filtering enabled the episodic queue models to achieve similar accuracy and trial completion rates to the persistent models, with approximately half the mean decision time. All models completed three decisions in nearly every trial, with a slight decrease in completion rate for the fourth decision.

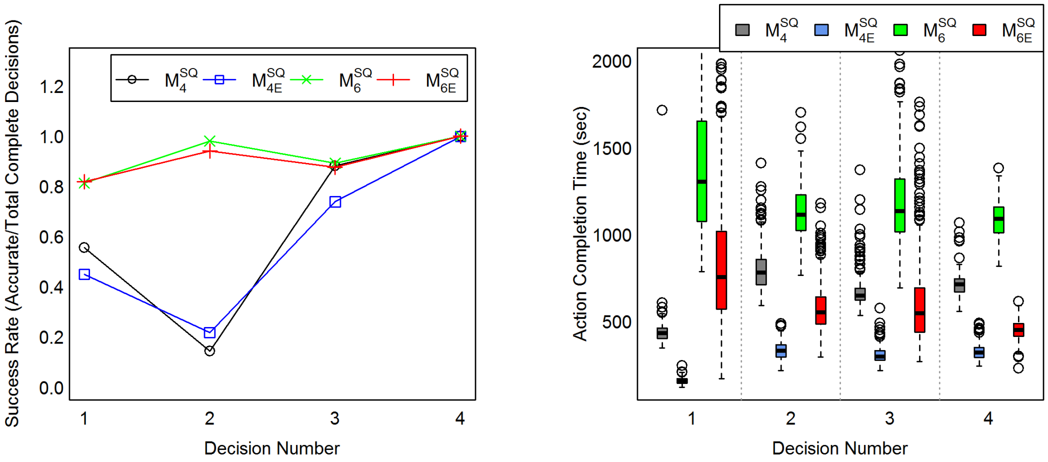

(a) Success Rate (b) Action Completion Time

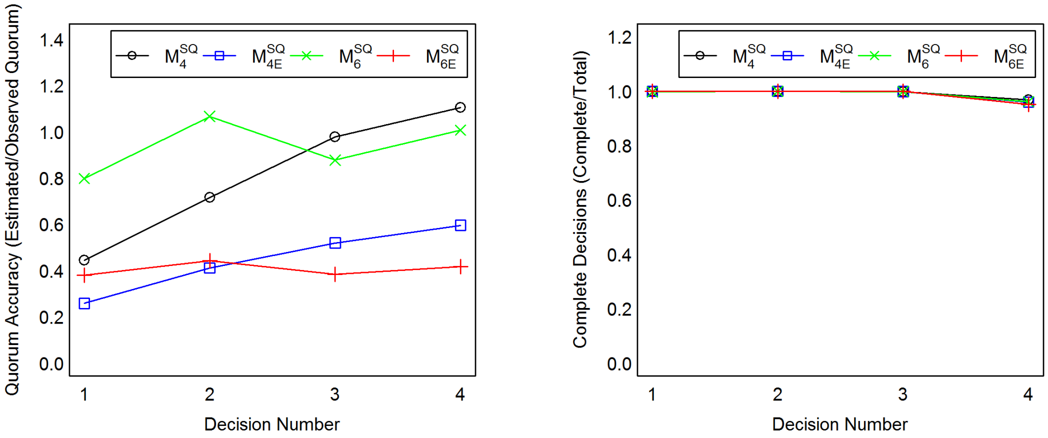

(c) Quorum Accuracy (d) Trial Completion Rate

Figure III.17: Strong Quorum Filtering (*SQ*): The models are compared by success rate (Fig. a), action completion time (Fig. b), quorum accuracy (Fig. c), and trial completion rate (Fig. d). Note: Fig. (b) cuts off the slowest 328 trials (3% of the total trials) with Action Completion Times exceeding 2000s (M_4-26,M_{4E}-1,M_6-261,M_{6E}-40) in order to facilitate viewing the primary results.

III.3.4.6 Discussion

The results indicated that the persistent queue models (M_4^* and M_6^*) were resilient to quorum detection noise, where the superscript * indicates all evaluated versions of the models. The episodic queue models (M_{4E}^* and M_{6E}^*), were significantly less resilient and did not complete trials without quorum noise filtering. Noise mitigation dramatically improved the ability of all models to complete multiple sequential decisions. Additionally, the results reinforced the results from the static site value experiment, as M_6^* and M_{6E}^* were more accurate, but slower than M_4^* and M_{4E}^*, respectively. Both of the initial hypotheses were partially supported.

Hypothesis $H_{4.3}$, which predicted higher trial completion for the persistent queue models (M_4^* and M_6^*)

was supported by the UQ, DI, and WQ scenarios, but not for the SQ scenario. The persistent queue models achieved high quorum accuracy for each decision, even in the unfiltered scenarios with normal (UQ) and delayed (DI) interaction periods. The episodic queue models were unable to make sequential decisions until queue filtering mechanisms were introduced. Greatly reducing the duplicate messages in the SQ scenario enabled the episodic queue models to achieve the same success and trial completion rates as the persistent models, but in significantly less time. A comparison of action completion times between of the episodic models in the unfiltered scenarios (UQ and DI) and the strong filtered scenario (SQ) showed that the episodic models required more time to complete actions in the strong filtering scenario. This observation exposes a weakness in the model when applied to situations that require the collective to move decision making hubs following decisions. Agents arriving at a new location support the same action execution. The absence of disagreeing agents at the new decision making hub location causes the agents' completed state quorum detection mechanism to serve as a timer that is regulated by the rate the queue is filled and the agents' quorum detecting threshold. Strong filtering slowed the rate that the queues received messages after agents arrived at the new decision making hub location; thus, preventing an early transition from the completed state the uncommitted state at the start of the collective's next decision.

The prior observation explains why the second hypothesis, $H_{4.4}$, which stated noise reduction techniques will improve the speed and completion rates of the models, is also only partially supported. The weak filtering queue scenario, WQ, resulted in all models having improved completion rates, but longer action completion times. Filtering queue messages by a small number of message identifiers reduced the computational requirements for the individual robots, but increased the probability that individual robots drop original messages with identifiers matching messages already in their queue. This design choice limited the pool of message identifiers to the queue length, k. After the agents transitioned to the new decision making hub, the weak queue filtering mechanism slowed the agents' transition from the completed state to the new decision making state enough to greatly improve the completion rate of the episodic models, but not enough to make the models reliable.

Split decisions are a significant problem that impede the reliable deployment of robotic collectives in situations requiring the collective to move their decision making hubs. Even though all models achieved high completion rates when using strong queue filtering, the remaining 3-5% failure rate is unacceptable for long sequences of collective action selections. A trivial solution assumes agents know the collective's population size and delay transitioning to a new decision until the agents estimate that an acceptable portion of the collective has arrived at the new location. This undesirable solution eliminates both the algorithm's scalability to different population sizes and prevents graceful degradation in the event of losing agents. One can either explicitly use timers based on expected action completion times to regulate transitions between decisions, or lengthen agent quorum detection queues. Increasing the queue length is predicted to further reduce, but not eliminate, the probability of the collective splitting. The splitting problem is less of an issue when the collective is not required to change decision making hub locations, since the completed agents will be in contact with agents that have not completed the desired task and be better able to estimate when the population is ready to transition, as evaluated in prior work (Parker and Zhang, 2010).

The sequential action selection experiment directly addressed research question R_4 by evaluating the presented models in scenarios requiring multiple sequential decisions. The persistent queue models (M_4^* and M_6^*) solved multiple collective action selection problems in sequence, even when the collective's individual robots were unable to filter quorum messages. The addition of weak (WQ) and strong (SQ) filtering enabled all evaluated models to consistently complete four action selections in sequence.

III.4 Summary

Novel implicit and explicit collective action selection models were introduced and evaluated in a series of four simulation experiments. The implicit action selection model extended an existing collective best-of-*n* decision making model in order reduce the influence of negative environmental bias during decision making. Interaction delays attempted to reduce the influence of disparate discovery times and travel times between sites in the collective's environment. Interaction frequency modulation adjusted individual interaction rates in order to compensate for distance between sites and the collective's decision making hub. The explicit action selection model enabled the collective to choose and execute actions in sequence. The presented set of site selection experiments evaluated each model's effectiveness. A summary of the research questions and evaluated hypotheses is provided in Table III.11.

Two experiments addressed research questions R_1-R_3 for collective site selection. The first experiment assessed the implicit action selection model for site selection with statically valued sites located at different distances to the collective's decision making hub. The interaction delay extension significantly improved the collective's decision accuracy, when the collective was required to choose sites based on their value, despite existing environmental bias due to site distance, which supported hypothesis H_1. The experimental results partially supported H_2. The interaction frequency modulation, which increased agent interaction attempts with the distance between the agent's supported site and the collective's decision making hub, did not improve decision accuracy, but did reduce decision time. The dynamic site values experiment addressed research question R_3 by evaluating the extensions for site selection subject to environmental bias and dynamically changing site values. The experiment supported hypothesis H_3 by demonstrating that the bias reducing model enabled consistent response to changes in site value during the decision making process. The response time for the original model varied significantly depending upon the location of the collective's considered options.

Table III.11: Summary of Hypotheses.

Question	Hypothesis		Experiment	Supported
R_1	H_1	Interaction delays will improve collective decision accuracy over the original model.	Static Site Value	Yes
R_2	H_2	Interaction frequency modulation will improve decision time and accuracy in the presence of environmental bias.	Static Site Value	Partially
R_3	H_3	Bias reduction causes uniform response to dynamic site value changes.	Dynamic Site Value	Yes
R_4	$H_{4.1}$	Episodic queuing techniques are faster and less accurate than persistent queuing techniques.	Quorum Detection	Partially
	$H_{4.2}$	Unanimous quorum thresholds cause slower, more accurate decisions.	Quorum Detection	Partially
	$H_{4.3}$	Persistent queuing techniques are more reliable than episodic queuing techniques in sequential decisions.	Sequential Action Selection	Yes
	$H_{4.4}$	Quorum noise reduction improves model speed and completion rate.	Sequential Action Selection	Partially

The quorum detection and sequential decision experiments addressed research question R_4 by evaluating the presented explicit collective action selection model for situations in which the collective chose and executed actions associated with a specific site. The quorum detection experiment required the collective to commit to a single site and stop deliberation. Majority and unanimous quorum thresholds were evaluated with persistent and episodic queuing techniques. The evaluated models consistently committed to the higher valued of the available sites, detected each decision, and stopped deliberation. The use of the persistent queue techniques caused slow decisions with high accuracy during quorum detection, while the episodic queue technique enabled faster quorum detection, but decision accuracy was close to the other models, which partially supported hypothesis $H_{4.1}$. The experiment's second hypothesis, $H_{4.2}$, was also partially supported by the experiment since unanimous quorum threshold models made slower decisions than the majority quorum threshold models, but both achieved similar accuracies. Only unanimous quorum thresholds were shown to prevent rare split decisions resulting from the collective committing to multiple sites simultaneously in the evaluated scenarios.

The sequential action selection experiment examined four versions of the explicit action selection model in scenarios requiring the collective to make a sequence of site selections. The models were distinguished by the presence or absence of the bias reducing extensions and the use of either persistent or episodic queues. Further, duplicate quorum detection messages were mitigated using increased delays between agent interactions and either weak or strong queue filtering. Hypothesis $H_{4.3}$ was supported, which demonstrated that the persistent queue models were able to complete a majority of the decision sequences, while episodic queues were only able to complete decision sequences with the addition of quorum noise filtering mechanisms. The introduction of quorum noise filtering did improve model completion rates, but not model speeds, which partially supports $H_{4.4}$. Split decisions caused the collective to fail to complete a sequence of decisions, which resulted from individual agents being left behind by the collective during movement to a new decision making hub.

The experimental results indicate that the implicit and explicit collective action selection models presented in this chapter provide reasonable approaches to two existing problems in collective decision making. First, most existing collective decision making models are inapplicable to scenarios subject to negative bias, since the collective's local environment alters the collective's decision outcome and speed. While the presented mechanisms are only applicable to environmental bias that causes disparate discovery and travel times between available options, the results in this chapter indicate that the presented models are resilient to different static and dynamic situations. The second problem is that of explicit action selection. Decoupling the actions a collective uses to assess available options from the actions the collective executes after a decision enables the collective to make decisions about what future action it will take. Prior work (see Chapter II.2.2) evaluated collective behavior transitions, but did not require the collective to choose its action prior to implementation. The results supported the presented explicit action selection model's ability to make and act on decisions; however, executing actions that move the collective's decision making hub were shown to split the collective in an unacceptable number of decisions. The use of high quorum thresholds and quorum noise mitigation reduced the probability of splitting, but did not eliminate the problem. Despite the discussed limitations of the models, each model offers the ability for humans and collectives to share decision making responsibilities in future human-collective teams. The use of explicit action selection in human-collective teams provides the ability for the collective to offer possible future action options to the human, which is a capability that is largely absent from the human-swarm interaction research.

Chapter IV

Human-Collective Team Action Selection

A novel approach to human-collective interaction is presented that also provides a framework for human supervision and control over explicit collective action selection (see Chapter III). The approach considers a subset of the collective autonomy, control, and transparency considerations introduced in Chapter II.3. The approach shares decision making between the human and the collective (Crandall et al., 2017; Coppin and Legras, 2012), while enabling the collective to operate at either a high or low level of autonomy, depending upon the human's choice of control and understanding of the collective. This chapter focuses on the last two research questions presented in Chapter I:

R_5**:** What mechanisms permit efficient human control over robotic collective action selection?

R_6**:** What visualization techniques improve human-collective interaction during collective action selection?

The chapter begins with a description of the interaction mechanisms and visualization techniques featured in the presented approach. A detailed description of the developed interface is then presented. Finally, a human subjects evaluation was conducted in order to analyze the interface's ability to provide the human with transparency into the collective's actions as well as the human's ability to influence the collective decision making algorithms (see Chapter III) in multiple-collective, multiple target scenarios.

IV.1 Human-Collective Interface

Requirements for the interface design focused on enabling a human to rapidly estimate the collectives' action selection processes, determine appropriate controls to use, and implement the controls. After these requirements are explained, the interface's simulated situation, control mechanisms, and visualization techniques are described.

IV.1.1 Interface Requirements

Three design requirements related to collective autonomy, control, and transparency (see Chapter II.3) guided the interface's development. First, the interface must enable the human to reduce or increase the collectives' levels of autonomy throughout the explicit collective action selection process. The collectives initially operate at high levels of autonomy and independently discover options, analyze options, select options and execute actions associated with each chosen option (see Table II.2). Using a set of control mechanisms, the human must be able to reduce the collective's autonomy in each of these decision making functions, as needed. Second, human interaction must be limited to the decision making hub. This requirement avoids assuming global communication capability, which eliminates from consideration existing control mechanisms, such as parameter setting and environmental controls (see Chapter II.3.3). Behavior selection controls, which enable the human to transition robots between internal states, are used as a means of simultaneously altering the collectives' levels of autonomy and guiding decision outcomes. Finally, the interface must enable the human to rapidly observe the collectives' decision making processes. Interaction with multiple collectives requires that the human be able to quickly determine if interventions into the collectives' decision making processes are needed and which controls to employ.

IV.1.2 Detailed Situation

The purpose of the developed interface was to enable human interaction with multiple collectives in an expanded version of the sequential action selection experiment described in Chapter III.3.4. A team comprised of a single human and four robotic collectives operated within an approximately 2 km^2 urban environment. Each collective consisted of a group of 200 simulated Unmanned Aerial Vehicles (UAVs) that were limited to a 500m exploration radius from the collective's central decision making hub. Robots behaved according to the individual model, as described in Chapter III.2.3.1, which enabled the collective to perform explicit action selection within its coverage area.

The team completed a series of targeting selection decisions in which each collective was required to choose the best target within its 500m search range, move its decision making hub to the chosen target, and initiate the collective's next target selection decision. Targets represented locations that the collectives assessed for future occupation. The targets were located in the environment and were invisible until discovered. The collectives' individual robots discovered a target and performed noisy (10%) assessments of the target's value when the robots were within sensory range of the target. After an individual robot discovered a target and returned to the collective's decision making hub, the target was added to the map and visible to the human. Once a target was occupied by a collective, the target was removed from the map and no longer discoverable. Collectives were able to find, evaluate, and occupy any target on the map that was within their search ranges, which made it possible for multiple collectives to consider the same target during a decision. After a collective moved, its next decision included all targets in range from the collective's new location.

Interactions between the individual robots and between the robots and the human were restricted to the collectives' decision making hubs. The human was assumed to have access to intermittent satellite imagery (simulated using the Google API) and relied on communication with robots within the collectives' decision making hub to persistently sense the environment. The human monitored the collectives' decision making processes and was able to issue requests to the appropriate collective in order to ensure each collective chose and moved to the highest valued target within range of the collective's robots. Additionally, the human was instructed to prevent multiple collectives from choosing and occupying the same target site. Rarely, two collectives moved to the same location (0.6% of all decisions), which caused the second collective to reach the target to automatically return to its previous position and initiated a new target selection decision.

The human-collective interaction assumed a larger set of individual robot capabilities than those used in the algorithm based evaluations (see Chapters III.3.1-III.3.4), which based individual robot movement and communication capabilities on the e-Puck (Mondada et al., 2009) robotic platform. Collective best-of-*n* studies typically use simple robots, such as Kilobots and e-Puck robots (Rubenstein et al., 2012; Mondada et al., 2009; Valentini et al., 2015), that have been optimized for laboratory research. Both Kilobots and e-Puck 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 ($<$ 25cm) peer to peer communication. Each capability is central to many biologically-inspired collective decision making strategies (Rubenstein et al., 2012; Mondada et al., 2009; Haque et al., 2016), but are difficult to recreate in UAV collectives. UAV peer-to-peer communication is complicated by the UAVs' high mobility (Sahingoz, 2014; Majumdar, 2015). Many commercial rotor wing UAVs are capable of speeds of 60km/h, but have limited flight times of approximately 25 minutes (e.g., Marketer, 2017). 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 30m peer-to-peer transmissions (e.g., optical transmissions (Kaadan et al., 2014)). A simplifying assumption was that robots moved at constant speed, which is similar to assumptions made in the algorithms evaluations (Chapter III.3), but notably deviates from practical UAV movement.

IV.1.3 Interface Visualization

The interface visualization needed to enable the human to rapidly estimate the state of the collectives' action selection processes, based on information received from individual robots within the collectives' decision making hubs. A screenshot of the developed interface is provided in Fig. IV.1. The interface was divided into four subareas. A visual representation of the four collectives' hubs and the respective discovered targets were geo-located on the map. The human used the Collective Requests area to change behavior states of agents within a collective's decision making hub using three requests described in Chapter IV.1.4. The Collective Assignments area provided a summary of active and inactive requests sent to the respective collectives. Finally, the System Messages area provided text alerts of the collectives' activities (e.g., finding targets) and feedback if the human made a mistake while issuing a request (Chapter IV.1.4).

Four collectives and sixteen targets are shown on the map in Fig. IV.1. The collectives were represented by white boxes identified by a Roman numeral, while the targets were represented by the green and blue boxes identified using an integer value. Left clicking on a hub or a target selected it as the designated object

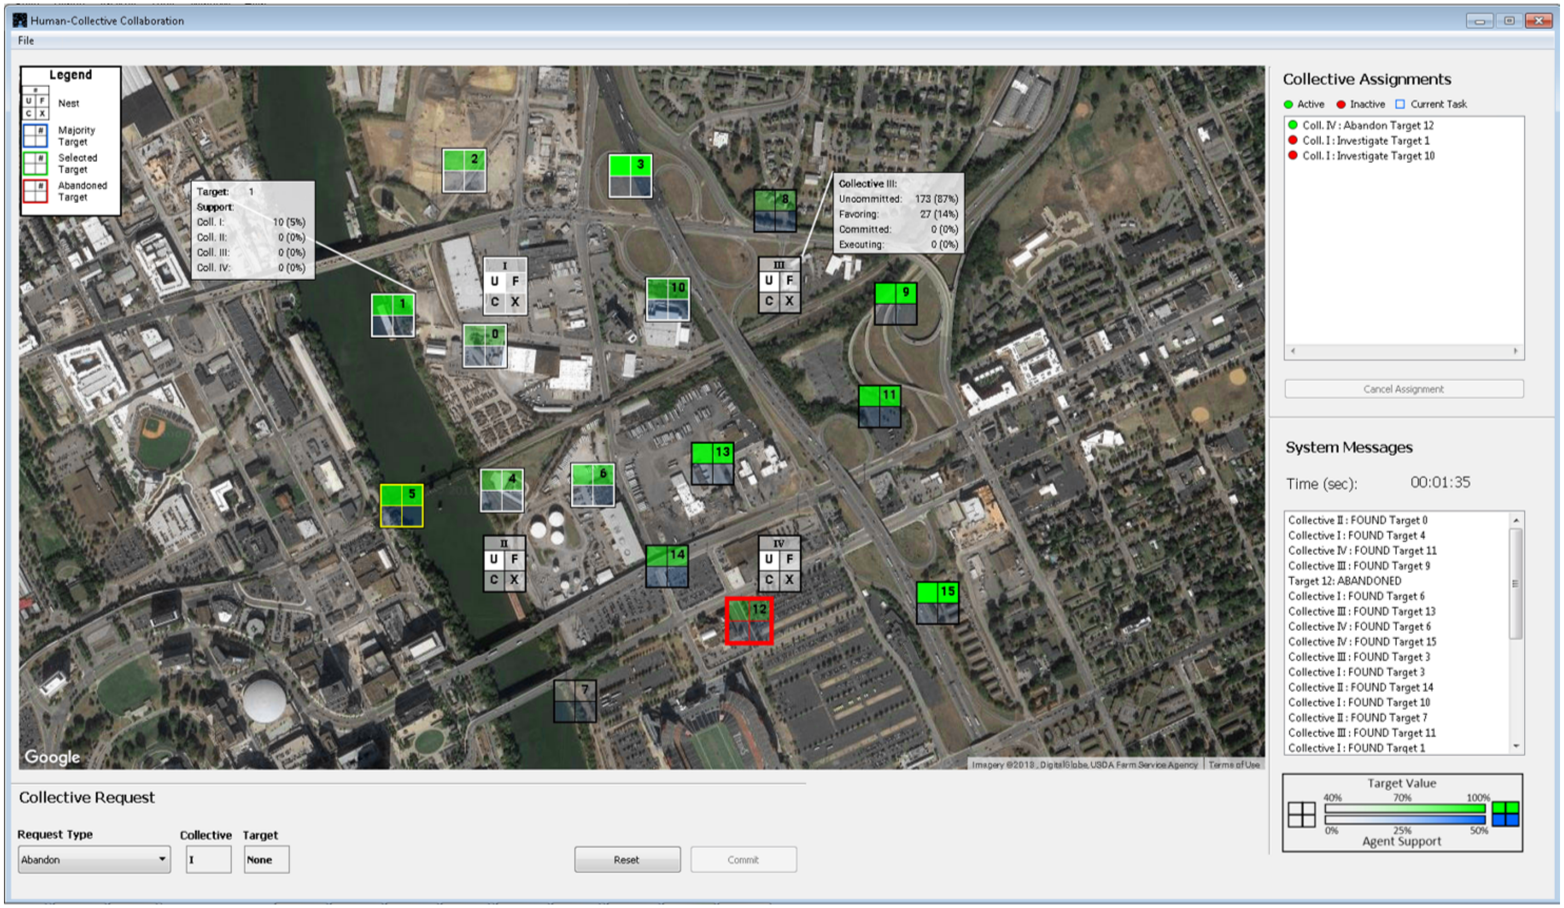

Target: 1
Support:
Coll. I: 10 (5%)
Coll. II: 0 (0%)
Coll. III: 0 (0%)
Coll. IV: 0 (0%)

Collective III:
Uncommitted: 173 (87%)
Favoring: 27 (14%)
Committed: 0 (0%)
Executing: 0 (0%)

Collective Assignments
- ● Coll. IV : Abandon Target 12
- ● Coll. I : Investigate Target 1
- ● Coll. I : Investigate Target 10

System Messages

Time (sec): 00:01:35

Collective II : FOUND Target 0
Collective I : FOUND Target 4
Collective IV : FOUND Target 11
Collective III : FOUND Target 9
Target 12: ABANDONED
Collective I : FOUND Target 6
Collective III : FOUND Target 13
Collective IV : FOUND Target 6
Collective IV : FOUND Target 15
Collective III : FOUND Target 3
Collective I : FOUND Target 3
Collective II : FOUND Target 14
Collective I : FOUND Target 10
Collective II : FOUND Target 7
Collective III : FOUND Target 11
Collective I : FOUND Target 1

Request Type: Abandon
Collective: I
Target: None

Figure IV.1: The human-collective interface mid-way through a trial 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 have indicates the support of the highest supporting collective (blue).

of the human's collective request (see Chapter IV.1.4). Left clicking on a hub also highlighted the discovered targets that the collective supported (with a white outline) and targets that were in range of the collective, but were not being supported (with a yellow outline). Right clicking on targets revealed an estimate of each collectives' support for the target, as shown with Target 1. Right clicking on a hub revealed a detail flag, as shown with Collective III, with an estimate of the number of robots in one of four states: *uncommitted* (U), *favoring* a target (F), *committed* to a target (C), or *executing* a move to a target (X).

Each collective's hub symbol contained four lettered quadrants representing each of the four robot states. Robots in the uncommitted state, U, either explored the environment looking for targets, or were within the decision making hub and available to be recruited by other robots. The number of robots that favored a target, F, were either periodically reassessing that target's value, or interacting with other agents within the decision making hub. The letter C represented the agents that had committed to executing a movement to a particular target after either detecting a quorum of support for that target or interacting with another committed robot. Finally, robots that were executing a movement, X, traveled from one decision making hub location to the target location chosen by the collective. The *initiating* and *complete* states of the individual behavior model (see Chapter III.2.3.1) were not presented directly to the human, but were aggregated into the *committed* and *executing* states, respectively. The opacity of each lettered quadrant indicated the relative percentage of agents in a particular state. Collective I in the figure is predominantly uncommitted (the U quadrant is more opaque and brighter than the other three quadrants).

Targets were represented by the blue and green boxes. A target's value estimate was more opaque and brighter green when it was higher. Newly discovered targets were not assigned a value by the interface until the interface received more than one report of the target's value. Early in the decision process, target values fluctuated as a result of the agent's noisy estimates of target value. The bottom half of each target symbol provided an indication of the agent support for that target with more opaque, brighter blue representing a higher level of support. When multiple collectives supported the same target, the blue color represented the collective with the highest level of support. The target detail flag (available by right clicking on the target) provided an estimate of support for each collective that supported the target. A centrally located target that was supported by all four collectives, for example, displayed each collective's support on its detail flag.

Finally, the interface attempted to alert the human to a collective's expected commitment and execution of movement to a chosen target. Collectives committed to targets when individual robots detected a majority of support for the specific target. A commitment resulted in the collective's transition to an execution state; thus, the interface provided an indication when a collective was likely to select a particular target by highlighting the target in blue when the collective's support exceeded 30%. The alert was intended to allow the human to intervene and stop the collective's decision, using an *Abandon* request, which is described in Chapter IV.1.4.

If the human did not modify the collective's decision, the target was highlighted with a green outline as robots began to execute the move to the new location. Finally, once a majority of the robots had left the current hub for the target, which became the new decision making hub, the hub representation moved from its current position to the chosen target's location, which indicated the start of the next decision cycle for that collective. Targets that were discovered by the collective previously remained on the map and the interface continued displaying the targets' previous values. The collective's next decision cycle included any targets on the map that were within range of the collective's new hub.

IV.1.4 Collective Controls

The human had the ability to select from three different behavior selection controls in order to alter the collectives' states via the robots within the decision making hubs. The collectives responded to human input, while executing the collective best-of-n decision making process. Once a collective committed to a target, it did not respond to human input until the collective initiated its next decision cycle.

Investigate requested that ten robots within a decision making hub transition from an uncommitted (U) state to a favoring (F) state supporting the human's chosen target, which subsequently increased support for the target within the collective. The number of contacted robots was calculated at the start of the experiment and based on 5% of the total population (10 robots in the evaluated scenario). Issuing the Investigate request required the human to specify the collective hub and the target to be investigated. The human was only aware of targets that the collectives had discovered and did not have any other additional information (e.g., intelligence reports). Once the request was committed, a record of the request appeared in the Collective Assignments area. The assignment remained Active (green), while the interface attempted to command agents within the decision making hub to make the transition. Robots that received the command responded with an acknowledgment. The assignment became Inactive (red), when the interface received acknowledgments from ten agents. Investigate requests for targets outside the collective's search range were ignored and an error message was displayed in the System Messages area.

The *Abandon* request reduced the collective's support for a specific target by transitioning robots in the hub that favored (F) a specific target to an uncommitted state (U). Issuing the Abandon request required the human to specify the collective hub and the target to be abandoned. Abandon requests remained active throughout the decision cycle, which means that the interface persistently attempted to transition all robots in the hub that supported the specified target to the uncommitted state. Unlike the Investigate request, the human had the ability to cancel an Abandon request (using the "Cancel Assignment" button). Abandon requests for newly discovered targets that had yet to be assigned a value (or green color) were ignored and an error message was displayed in the System Messages area.

Decide transitioned agents that already supported a target to commit (C) to that target, which resulted in the same effect as agents detecting a quorum of support. The decide request remained active, while the interface attempted to command ten robots to commit; however, once the first commanded robot committed, it immediately began communicating the commitment to other robots within the collective's decision making hub. The decide requests were limited to targets that had at least 30% of the collective's support or higher. Decide requests for targets with less than 30% support from the selected collective were ignored and resulted in an error message.

These requests enabled the human to influence the algorithm's behavior by altering, accelerating, or delaying the collectives' target selections and action executions. A notable absence from the presented approach was the ability for the human to suggest locations for investigation (e.g., Crandall et al., 2017) and to receive additional environmental information from the collective. These and several additional interactions require further investigation and are discussed in Chapter V.

IV.2 Human-Collective Site Selection Experiment

The purpose of the human-collective site selection experiment was to evaluate the presented human-collective interaction approach in a simple scenario that required multiple site selections that were subject to environmental bias. The goal of the experiment was twofold. First, different collective action selection behaviors were evaluated for effectiveness in a human-collective decision making scenario. Second, the presented human-collective interaction approach was evaluated in order to determine its ability to enable the participant to influence, understand, and control robotic collective action selection. Four hypotheses were tested. Hypothesis $H_{5.1}$ was that the presented behavior selection control mechanisms enable users to control collective action selection in robotic collectives. Hypothesis $H_{5.2}$ predicted that the bias reducing models will reduce users' workload by requiring less human input to make accurate targeting decisions. Hypothesis $H_{5.3}$ predicted that human input improves the accuracy and decision times for the collective action selection models. Finally, hypothesis H_6 was that the presented visualization techniques were sufficient to enable users to interact with the collective and understand collective behavior.

IV.2.1 Experimental Design

The mixed design experiment consisted of three twenty minute targeting trials that paired a participant with one of three collective action selection behaviors. Each trial required participants to supervise and influence between 12 and 16 target selection decisions made by four robotic collectives simultaneously in a simulated environment. The variance in the number of target selection decisions represents the minimum and maximum number of target selection decisions that the participants completed within the time limit. Target selection decisions were made when the collective, with or without human user input, achieved a majority of support for a target, moved its decision making hub to the chosen location, and initiated the next targeting decision. Collectives were required to make two target selection decisions in sequence. The target selection decisions were distinguished by difficulty. Hard decisions occurred when the optimal target was located further from the collective decision making hub than other closer, mediocre targets. Easy decisions occurred when the optimal and mediocre targets were equidistant from the collective's decision making hub, or the optimal targets were located closer to the decision making hub than the mediocre targets.

The evaluated models were the primary independent variable and are summarized in Table IV.1. The first two models, M_{4E}^{SQ} and M_{6E}^{SQ}, were the strong quorum filtering models used in the sequential action selection experiment (see Chapter III.3.4). These particular models were selected, because they successfully and more quickly completed sequences of decisions in the sequential action selection experiment, than the other evaluated models. Model M_{4E}^{SQ} performed basic explicit action selection (see Chapter III.2.3). The M_{6E}^{SQ} model extended M_{4E}^{SQ} with the bias reduction mechanisms, described in Chapter III.2.2, in order to compensate for sites that were located at different distances from the collective's decision making hub. The final evaluated model, M_7, explored for evaluated targets in the environment, but did not use a collective best-of-*n* decision making strategy. M_7 was included in order to establish participants' baseline performance when using a low autonomy collective for comparison with participant's performance with the higher autonomy collective decision making algorithms. The secondary independent variables included the locations of the targets with respect to each collective's decision making hub, the targets' values, and the decision difficulty (e.g., easy or hard). All target values were static.

At the start of each trial, the four collectives were positioned 600m apart in a square formation, centered within the environment, which purposely overlapped the collective 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 $\{0.7, 0.8, 0.9, 1.0\}$. Target values were assigned to targets based on initial decision problem difficulty. Easy initial problems were arranged with the optimal site value (1.0) closest to the collective decision making hub (e.g., at distance of 150m), while hard initial problems were arranged with the optimal site value (1.0) furthest from the decision making hub (e.g., at a distance of 350m).

Table IV.1: Evaluated Models

Model	Description
M_{4E}^{SQ}	Action selection with Episodic, Strong Filtered Queue
M_{6E}^{SQ}	M_{4E}^{SQ} with Bias Reduction
M_7	Exploration Only Model

Three trials were conducted, where a single trial corresponded with one of the collective action selection models. The first two trials were assigned, such that half the participants completed the evaluation with M_{4E}^{SQ} followed by M_{6E}^{SQ}, and the other half completed the M_{6E}^{SQ} model first. The baseline model M_7 was expected to benefit from learning effects related to the problem, interface, and collective behavior when compared to the primary models; thus, it was placed last in the sequence.

A number of objective metrics were gathered. The collectives' actions and decision making states were recorded automatically by the interface software, as were the participants' request and observation activities conducted via the interface. Participant requests included attempts to influence a collective's decision making process using the previously described *Investigate*, *Abandon*, and *Decide* requests. The recorded participant observation activities included opening detail flags and examining targets supported by the collective. Finally, the participants were required to respond to twelve real-time situational awareness probes during each trial, at a rate of approximately one question per minute.

Situational Awareness (SA) is a participant's understanding of the current and future state of the environment (Endsley, 1995). Probing questions have been used in prior research in order to determine a participant's situational awareness in critical tasks (Humphrey, 2009). The probes were designed to determine the participant's situational awareness according to Endsley's (1995) three levels of situational awareness: perception, comprehension, and projection. Perception (Level 1) questions determined the participants' 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 participants' 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 participants' ability to estimate the system's future state based on their perception and comprehension of the current state (e.g., "Which collective will make the next selection?"). Participants were asked four Level 1, five Level 2, and three Level 3 SA probes for each model. The SA probes consisted of question templates that asked similar questions, but due to human influence on the collectives, specific collectives and targets in the questions varied. 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.

The objective metrics that were used to compare the performance of the human-collective teams across the models included: Success Rate (SR), Decision Time (DT), Request Frequency (RF), and Intervention Frequency (IV). Success rate was the ratio of the correct decisions, made by either the human or the collective, to the total number of decisions made by the human-collective team. A correct decision occurred when the collective executed an action associated with the highest valued site found by the collective (see Definition 7). Decision Time was the time from the start of a collective decision to the successful completion of the collective's chosen action, with or without human input. Request Frequency was the number of requests, per minute, issued by the participants during the decisions. Each participant input was associated with a specific collective, target, and time. Requests for each individual collective were associated with specific decisions, if the requests were sent to the collective between the start and finish of a decision problem. The request frequency was the ratio of the number of requests each collective received to each collective's decision times. Finally, the intervention rate was the ratio of the participants' interventions to the total number of decisions made by the human-collective team. An intervention was defined as an Abandonment request issued to a collective after the collective had achieved 20% or higher support for a particular target. This action indicated that the participant was required to override the collective's independent decision making process.

Additional objective metrics were used to compare the participants' observation activities and situational awareness between trials. Observation Frequency (OF) was the number of observation actions, per minute made by the participants during decisions. Observation actions included opening detail flags and examining the collectives' targets. The correct Situational Awareness probe response ratio (SA) was the ratio of correctly answered situational awareness probe questions to the total number of questions asked within a trial. Probe responses were correct if the participant answered the question correctly. Probe responses were incorrect either due to an incorrect response, or if the participant failed to respond after one minute. Questions were read clearly two times in succession, to ensure the participant heard the question, but were not read a third time. The term SA_O corresponds to the overall correct probe responses, while SA_1, SA_2, and SA_3 were used to refer to level one, two, or three situational awareness probes, respectively.

Several sources of subjective data were collected. Prior to beginning the experiment, participants completed a demographic questionnaire and Vandenberg and Kuse's (1978) Mental Rotations Test (MRT). The demographic questionnaire gathered basic information related to sex, age, education level, and familiarity with office workstation systems. The MRT is a ten minute pencil and paper test that assessed each participant's spatial abilities and is distributed by Peters et al. (1995). After each trial, the participants completed the NASA Task Load Index (NASA-TLX), which determined perceived workload; the 3-D Situation Awareness Rating Technique (Selcon et al., 1991); and a post-trial questionnaire focused on the perceived performance and the responsiveness of the collective during the trial. Finally, the participants completed a post-experimental questionnaire, which requested participants to compare the three algorithms based on overall performance. The post-experimental questionnaire required rank ordering the models based on the model's responsiveness to the participants' requests, the participants' ability to make accurate decisions with the collective, and the participants' perceived understanding of the model.

Each participant began the experiment by completing the informed consent paperwork, the demographic questionnaire, and the MRT-A. Once these items were completed, participants received a scripted introduction to the experiment and the simulator (5 minutes). Participants were permitted to ask questions during the introduction and prior to the initiation of the initial trial. Each participant conducted a total of three trials (one trial for each evaluated model). Prior to each trial, the participants conducted 5 minute training sessions consisting of two collectives, with one collective required to solve an easy problem and one collective required to solve a hard problem. The collectives implemented the same action selection model as the collectives in the trial. Problem difficulty was not discussed with the participants. After the training session, the trial was conducted during which four collectives were supervised making simultaneous action selections. Trials consisted of two ten minute sections, distinguished by either hard or easy initial collective decisions. Following a collective's first target selection decision and the occupation of the collective's first chosen target site, the collective's next decision was either easy or hard depending upon the collective's location relative to the remaining targets. Each trial section ended when the participant (or the collectives) had completed eight target selection decision problems (two decisions per collective), or more than ten minutes had elapsed. Between trial sections, the collectives were reset to their initial locations, all targets from the previous part of the trial were removed, and the map was populated with targets from the the next trial section. The order of the hard and easy trial sections was counterbalanced across participants. Participants responded to the situational awareness probes at increments of approximately one probe per minute. At the end each trial, the participant completed the post-trial questionnaire, the NASA-TLX, and the 3-D SART. After completing all trials and associated post-trial data collection, the participant completed the post-experiment questionnaire.

Twenty-eight participants from the Vanderbilt University campus and surrounding area completed the experiment. The 15 females and 13 males were predominantly in the 18 to 30 year age range, although four participants were between the ages of 31 and 50. The participants had completed high school and most (13 participants) had either completed or were completing (11) an undergraduate degree. The MRT-A has a minimum score of 0 and a maximum possible score of 24. The participants' MRT-A scores resulted in a 10.9 mean score, with a standard deviation of ±5.5, a median score of 10, and minimum and maximum scores of 1 and 24, respectively. These results were virtually identical to the results obtained by Peters et al. (1995), who reported a 10.8 mean score, with a standard deviation of ±5 in a study comprised of 636 participants.

IV.2.2 Objective Results: Performance

The descriptive statistics for Success Rate (SR), Decision Time (DT), Request Frequency (RF), and Intervention Frequency (IV) are presented in Table IV.2. The highest success times, decision times, and request frequencies were observed for the M_7 model. The lowest success rates and fastest decision times occurred when using the M_{4E}^{SQ} model. The bias reducing model, M_{6E}^{SQ} enabled a success rate that approached the success rate for the M_7 model, but required less time and approximately half the requests. Intervention frequency was

Table IV.2: Descriptive Statistics by Model and Objective Performance Metric.

Model	Statistic	SR	DT	RF	IV
M_{4E}^{SQ}	Mean (SD)	0.79 (0.41)	3.24 (1.65)	0.87 (0.61)	0.38 (0.49)
	Median	1	2.75	0.79	0
	Min (Max)	0 (1)	1.16 (10.32)	0 (2.72)	0 (1)
M_{6E}^{SQ}	Mean (SD)	0.89 (0.32)	4.08 (1.38)	0.91 (0.57)	0.18 (0.38)
	Median	1	3.72	0.81	0
	Min (Max)	0 (1)	1.83 (9.94)	0 (2.92)	0 (1)
M_7	Mean (SD)	0.93 (0.25)	5.17 (2.16)	1.68 (0.76)	0.16 (0.37)
	Median	1	4.66	1.62	0
	Min (Max)	0 (1)	1.48 (13.25)	0.37 (4.15)	0 (1)

highest when using the M_{4E}^{SQ} model, which indicates that participants were required to alter the collective's behavior more than twice as often in order to achieve accurate results compared to the remaining models. Interestingly, the intervention frequencies for M_{6E}^{SQ} and M_7 were nearly identical, suggesting participants intervened in order to alter the M_{6E}^{SQ} collective's behavior at roughly the same rate as with M_7.

A Kruskal-Wallis test indicated significant effects for success rate, $\chi^2(2, N = 1258) = 37.726$, $p < 0.001$, decision time, $\chi^2(2, N = 1258) = 217.98$, $p < 0.001$, request frequency, $\chi^2(2, N = 1258) = 277.95$, $p < 0.001$ and intervention frequency $\chi^2(2, N = 1612) = 83.276$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for success rate between M_{4E}^{SQ} and M_{6E}^{SQ} ($p = 0.0458$) and between M_{4E}^{SQ} and M_7 ($p < 0.01$). Significant effects for success rate were not observed between M_{6E}^{SQ} and M_7. Significant effects for decision time were observed for all model comparisons ($p < 0.001$). Significant effects for request frequency was observed in both pairwise comparisons with M_7 ($p < 0.001$), but no significant effects were observed between M_{4E}^{SQ} and M_{6E}^{SQ}. Finally, pairwise comparisons revealed significant effects for intervention frequency between M_{4E}^{SQ} and both M_{6E}^{SQ} and M_7 ($p < 0.001$), but no significant effects were observed between M_{6E}^{SQ} and M_7.

A comparison of each model's performance by problem difficulty is provided in Fig. IV.2. Easier decisions, which were not subject to negative bias, resulted in high success rates across all models, as shown in Fig. IV.2a. Difficult problems were influenced by negative bias, which resulted in M_{4E}^{SQ} having the lowest success rates. M_7 achieved the highest success rate, although it was not significantly higher than M_{6E}^{SQ}'s success rate. Decision times, shown in Fig. IV.2b, increased for all models between the easier and more difficult problems, with model M_7 performing the slowest overall. Model M_{4E}^{SQ} was clearly the fastest model when making easy decisions, indicating that its decision process was accelerated when optimal sites were closer to the collective decision making hub. M_{6E}^{SQ} was slower than M_{4E}^{SQ} in easy problems, but both models' decision times were nearly identical during hard problems, even though M_{6E}^{SQ} was significantly more accurate. Request frequency was similar for both M_{4E}^{SQ} and M_{6E}^{SQ} in both problem difficulty levels, which was unexpected. Fur-

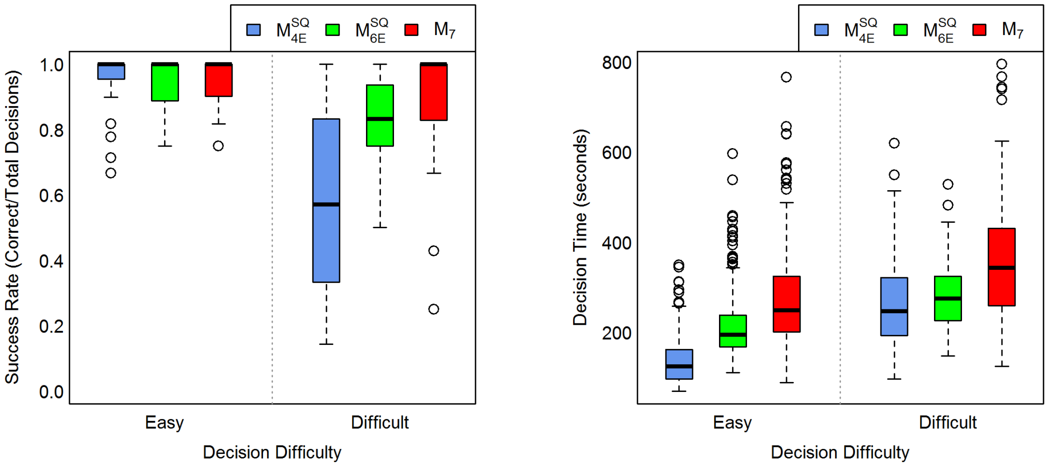

(a) Success Rate

(b) Decision Time

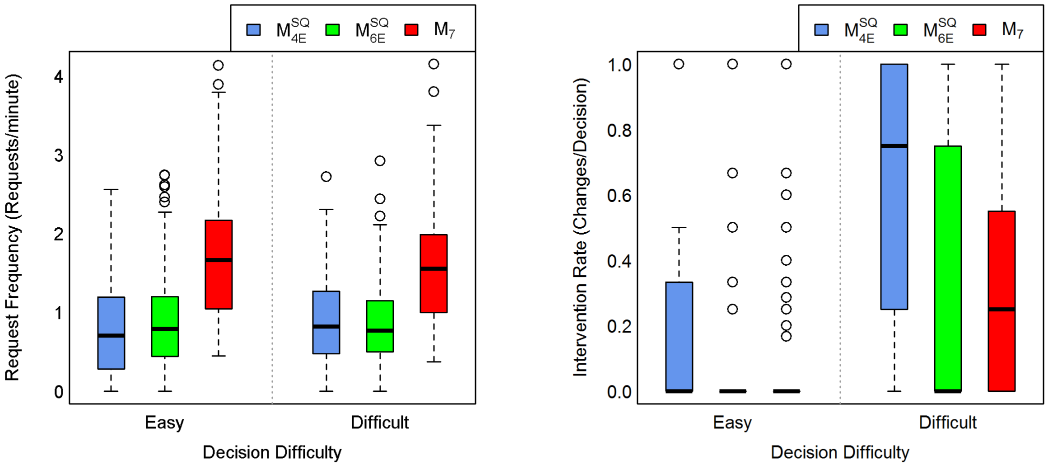

(c) Request Frequency

(d) Intervention Frequency

Figure IV.2: The performance of the evaluated collective action selection models in human-collective targeting trials. The models are compared based on success rate (a), decision time (b), request frequency (c), and intervention frequency (d).

ther analysis revealed that M^{SQ}_{4E}'s request frequency was weakly correlated with its success overall, $r = 0.216$, $p < 0.001$, and moderately correlated for difficult problems, $r = 0.439$, $p < 0.001$. No correlation was observed between request frequency and success rate M^{SQ}_{6E}, which indicates M^{SQ}_{4E} was more reliant on participant actions in order to make correct decisions. M_7 was expected to require the highest number of requests in order to complete a decision, since it did not include a collective best-of-n process. Finally, the intervention rates, shown in Fig. IV.2d indicate the rate, per decision, that participants changed a collective's behavior by abandoning a site the collective had identified and favored independently. Participants were more likely to be required to abandon sites supported by M^{SQ}_{4E} in order to make accurate decisions, when compared to the M^{SQ}_{6E} model. Difficult problems required participants to abandon M^{SQ}_{4E}'s selected sites for the largest num-

Table IV.3: Descriptive Statistics by Model for Request Types per Decision and Request Types per Minute.

		Number of Requests per Decision			Number Requests per Minute		
Model	Statistic	Investigate	Abandon	Decide	Investigate	Abandon	Decide
M_{4E}^{SQ}	Mean (SD)	1.81 (1.97)	0.72 (1.06)	0.33 (0.48)	0.55 (0.51)	0.19 (0.27)	0.11 (0.18)
	Median	1	0	0	0.45	0	0
	Min (Max)	0 (12)	0 (7)	0 (2)	0 (2.56)	0 (1.8)	0 (0.8)
M_{6E}^{SQ}	Mean (SD)	2.31 (2.09)	0.7 (1.06)	0.53 (0.51)	0.58 (0.5)	0.16 (0.25)	0.15 (0.16)
	Median	2	0	1	0.43	0	0.15
	Min (Max)	0 (13)	0 (6)	0 (2)	0 (2.63)	0 (1.2)	0 (0.55)
M_7	Mean (SD)	4.99 (2.46)	1.53 (1.69)	1.04 (0.25)	1.13 (0.74)	0.27 (0.28)	0.24 (0.12)
	Median	5	1	1	1.01	0.22	0.21
	Min (Max)	0 (14)	0 (10)	1 (4)	0 (3.54)	0 (1.11)	0.08 (1.15)

ber of decision problems. Corrections were necessary in some instances for the M_{6E}^{SQ} model during difficult decisions, but the model maintained a low median intervention rate. The median intervention rate for the baseline model, M_7, was higher than anticipated, suggesting that the participants supported one site early in the decision process and were required to abandon the site after finding another higher valued site.

The descriptive statistics for the Investigate, Abandon, and Decide requests per decision and per minute for each model are provided in Table IV.3. The requests used for the M_7 model provide an estimate of the number of requests participants used per decision in the absence of a collective best-of-*n* process. The results indicate that participants issued Investigate and Decide requests more often per decision for the M_{6E}^{SQ} model, than for the M_{4E}^{SQ} model. The number of Abandon requests used per decision was roughly the same. The prior comparison of decision times (see Fig. IV.2) revealed that the M_{6E}^{SQ} model was slower than the M_{4E}^{SQ} model during easy decision problems, but that the models had similar decision times for hard decision problems. The right side of Table IV.3 presents the use of the different requests by minute. The per minute use of the Investigate and Abandon requests was similar between models M_{6E}^{SQ} and M_{4E}^{SQ}. Decide requests were only used once per decision.

Kruskal-Wallis tests were used to determine significant effects for the number and type of request per decision (left side of Table IV.3). The test indicated significant effects for the Investigate requests per decision, χ^2 $(2, N = 1258) = 381.75$, $p < 0.001$, Abandon requests per decision, χ^2 $(2, N = 1258) = 73.654$, $p < 0.001$, and Decide requests per decision, χ^2 $(2, N = 1258) = 419.59$, $p < 0.001$. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test revealed significant effects for the per decision number of Investigate requests and Decide requests between all models ($p < 0.001$). The pairwise comparisons revealed significant effects for the number of Abandon requests per decision between M_7 and M_{4E}^{SQ} ($p < 0.001$) and between M_7 and M_{6E}^{SQ} ($p < 0.001$).

Analysis of the number of requests per minute (right side of Table IV.3) produced different results than those observed for requests per decision. A Kruskal-Wallis test indicated significant effects for the Investigate requests per minute, χ^2 $(2, N = 1258) = 197$, $p < 0.001$, the Abandon requests per minute, χ^2 $(2, N = 1258) = 43.713$, $p < 0.001$, and the Decide requests per minute, χ^2 $(2, N = 1258) = 191.42$, $p < 0.001$. The results of the pairwise comparisons using a Tukey and Kramer (Nemenyi) test did not identify any significant effects between models M_{4E}^{SQ} and M_{6E}^{SQ} for either the Investigate or Abandon requests per minute. The test revealed significant effects between both models and the baseline model (M_7) for Investigate requests per minute ($p < 0.001$) and Abandon requests per minute ($p < 0.001$). Finally, the test indicated significant effects in all pairwise comparisons for the number of Decide requests per minute ($p < 0.001$).

Further analysis reveals weak correlations between the request frequency and success of the M_{4E}^{SQ} model over all decisions for the Investigate request, $r = 0.139$, $p < 0.001$, and the Abandon request, $r = 0.1076$, $p = 0.0256$. When only hard problems were analyzed for M_{4E}^{SQ}, the correlations were moderately stronger between success rate and the Abandon frequency, $r = 0.3365$, $p < 0.001$, and between success rate and the Investigate frequency, $r = 0.2749$, $p < 0.001$. No correlations were observed for either requests or decision difficulties for the M_{6E}^{SQ} model. Additionally, correlations did not exist between request frequencies and success rates in the easy decision problems for either model. This additional analysis suggests that the participants' use of the Investigate and Abandon requests influenced model M_{4E}^{SQ}'s success rate, but was less important for success with model M_{6E}^{SQ}.

Although not presented in the primary data, the strong quorum filtering mechanism (Chapter III.3.4.5) greatly reduced the occurrence of split decisions when compared to prior results without quorum noise reduction (see Table III.7). During the human-collective trials 0.05% (6 of the $1,258$ targeting decisions) caused split decisions to occur. While this percentage is small, the trials occurred in an idealized situation that did not require more than two decisions to occur in sequence. Prolonged use of the collective action selection strategy in real-world systems will need to better address the prevention and recovery from collective splitting.

IV.2.3 Objective Results: Independent Collective Performance

Each of the trials used during the human subjects evaluation were also conducted with the collectives independent of the participants, in order to determine the effect of human interaction on the collectives' performance. The descriptive statistics for the independent collectives' performance according to Success Rate

Table IV.4: Independent Collective Performance Descriptive Statistics by Model and Metric.

Model	Statistic	SR	DT
C_{4E}	Mean (SD)	0.5 (0.5)	2.86 (1.28)
	Median	1	2.55
	Min (Max)	0 (1)	1.05 (13.48)
C_{6E}	Mean (SD)	0.72 (0.45)	5.21 (2.05)
	Median	1	4.76
	Min (Max)	0 (1)	1.65 (20.28)

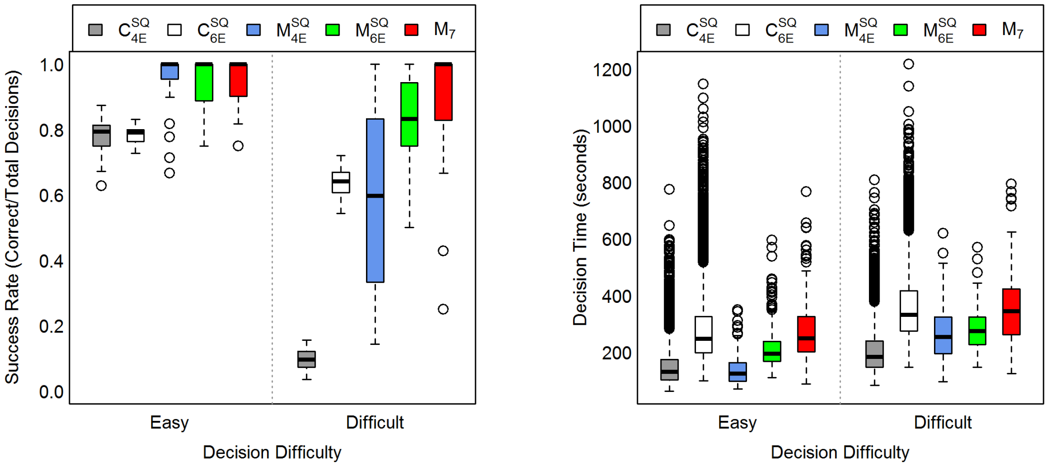

(a) Success Rate

(b) Decision Time

Figure IV.3: The independent collective performance compared to the human-collective performance, according to Success Rate (a) and decision time (b).

(SR) and Decision Time (DT) are presented in Table IV.4. The collective models C^{SQ}_{4E} and C^{SQ}_{6E} represent independent collectives implementing models M^{SQ}_{4E} and M^{SQ}_{6E}, respectively. Consistent with the human subjects evaluation, the bias reducing model, C^{SQ}_{6E}, achieved higher accuracy and slower decision times when compared to C^{SQ}_{4E}. A Kruskal-Wallis test indicated significant effects between the models for success rate, $\chi^2(1, N = 44784) = 2233.9$, $p < 0.001$ and decision time $\chi^2(1, N = 44784) = 18285$, $p < 0.001$.

The performance of the independent collectives is compared to the performance of the human-collective teams in Fig. IV.3. The independent collectives made accurate decisions in approximately 80% of the easy problems, as shown in Fig. IV.3a. During difficult decision problems, C^{SQ}_{4E} was unable to reliably make accurate decisions. Model C^{SQ}_{6E}, which incorporated the bias reducing mechanisms, achieved an independent success rate during difficult problems that was comparable to the success rate achieved by the human-collective team without bias reduction, indicated by the success rate of M^{SQ}_{4E}. Comparing the success rates of C^{SQ}_{4E} and C^{SQ}_{6E} to the human-collective success rates (M^{SQ}_{4E} and M^{SQ}_{6E}) indicates that the participant interactions improved collective success rates for both evaluated behaviors. Model C^{SQ}_{4E} made faster decisions than C^{SQ}_{6E} during easy and difficult problems, as shown in Fig. IV.3b. Human interaction with these behavior improved decision times during easy problems, as indicated by the decision times achieved by M^{SQ}_{4E} and M^{SQ}_{6E}. During difficult decision problems, participant interaction improved the decision times of the bias reduction model (M^{SQ}_{6E}), but increased the decision time of the non-bias reducing model (M^{SQ}_{4E}). The increased decision time of M^{SQ}_{4E} compared to the independent collective (C^{SQ}_{4E}) is indicative of the participants' need to intervene in order to achieve accurate decisions with the non-bias reducing model.

Table IV.5: Descriptive Statistics by Model and Objective Situational Awareness Metric.

Model	Statistic	OF	SA Probe	SA$_1$	SA$_2$	SA$_3$
M_{4E}^{SQ}	Mean (SD)	0.66 (0.68)	0.85 (0.19)	0.86 (0.26)	0.86 (0.26)	0.8 (0.37)
	Median	0.54	0.83	1	1	1
	Min (Max)	0 (3.85)	0 (1)	0 (1)	0 (1)	0 (1)
M_{6E}^{SQ}	Mean (SD)	0.64 (0.54)	0.9 (0.14)	0.91 (0.17)	0.88 (0.21)	0.9 (0.3)
	Median	0.54	1	1	1	1
	Min (Max)	0 (2.66)	0.33 (1)	0.5 (1)	0.5 (1)	0 (1)
M_7	Mean (SD)	0.57 (0.48)	0.94 (0.1)	0.94 (0.19)	0.91 (0.2)	0.97 (0.12)
	Median	0.48	1	1	1	1
	Min (Max)	0 (2.42)	0.67 (1)	0 (1)	0 (1)	0.5 (1)

IV.2.4 Objective Results: Situational Awareness

The descriptive statistics for the Observation Frequency (OF) and the correct Situational Awareness probe response ratio (SA) are presented in Table IV.5. The observation frequencies were similar for all models, although the M_7 model required slightly less frequent observations. The SA probe correct response ratio was high for all models, although participants were noticeably more likely to respond incorrectly to the SA probes when using model M_{4E}^{SQ} overall (SA_O), and for each situation awareness level (SA_1, SA_2, and SA_3).

A Kruskal-Wallis test did not identify any significant effects for observation frequency, or for SA probe response. The chosen objective measurements of the participants' situational awareness were largely indifferent to the models. The lack of significance in observation frequency is only somewhat unexpected, since the experimenter observed that participants tended to establish an observation routine quickly that did not appear to vary between trials. The slight differences in correct probing question responses between the models are not conclusive since there were no significant effects between models M_{4E}^{SQ} and M_{6E}^{SQ}. The higher response rate for the M_7 model suggests either that the participant benefited from controlling the collective, or by the third trial, the participant was more familiar with the interface in general.

IV.2.5 Subjective Results

This experiment considered several sources of subjective data. The first set of subjective results are the participants' reported workload (NASA-TLX) and situational awareness (3D SART). The NASA-TLX weighted score provided perceived workload where each of the workload components (e.g., mental demand, effort, and frustration) were weighted according to each component's relative value. The 3D SART score ($SART$) for situational awareness is 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 descriptive statistics are presented in Table IV.6. The participants' perceived workloads with each model were similar, although model M_{4E}^{SQ} was slightly higher than the other models. The participants' reported situational awareness (SART Score) was also similar across all models,

Table IV.6: Descriptive Statistics by Model and Subjective Metrics.

Model	Statistic	NASA-TLX	SART Score	SU	DAR	SAR
M_{4E}^{SQ}	Mean (SD)	58.31 (17.63)	6.05 (2.22)	5.74 (0.99)	4.98 (1.2)	5.28 (1.45)
	Median	60.67	6	6	5	6
	Min (Max)	9 (89.33)	1 (13)	3 (7)	1 (7)	2 (7)
M_{6E}^{SQ}	Mean (SD)	57.06 (16.47)	6.75 (2.23)	6.08 (0.89)	5.04 (1.19)	5.71 (1.06)
	Median	56.83	7	6	5	6
	Min (Max)	5.67 (83.33)	3 (13)	4 (7)	1 (6)	3 (7)
M_7	Mean (SD)	50.63 (17.56)	6.38 (1.99)	6.09 (0.96)	5.07 (1.41)	5.36 (1.3)
	Median	54.17	6	6	5	5
	Min (Max)	9.33 (80.33)	4 (11)	4 (7)	1 (7)	3 (7)

although situational awareness with model M_{6E}^{SQ} was slightly higher than the other models, due primarily to M_{6E}^{SQ}'s higher reported supply of attentional resources (*SAR*) when compared to models M_{4E}^{SQ} and M_7. The Kruskal-Wallis test did not identify any significant effects for workload, situational awareness, or any of the three components of the 3D SART score. The participants' 3D SART scores were consistent with the correct SA probe response ratios (see Table IV.5).

The post-trail questionnaires required Lickert scale responses, on a scale of 1 (worst) to 7 (best) regarding: collective responsiveness to requests (e.g., *Investigate*, *Abandon*, or *Decide*), the collective's independent target selection ability (Performance), and ease of understanding the collective (Understanding). The descriptive statistics are presented in Table IV.7. The post-trial questionnaire responses indicate that participants' perceived that the collectives responded similarly to each type of request. The results for the Investigate and Abandon commands for models M_{4E}^{SQ} and M_{6E}^{SQ} correspond to the objective results (see Table IV.3), which indicated that both commands were used equally. The M_{4E}^{SQ} model was perceived to be slightly more responsive to Decide requests than M_{6E}^{SQ}, although the request was used more frequently in M_{6E}^{SQ} (Table IV.3). Perceived model performance was observed to be slightly lower for M_{4E}^{SQ}, which corresponds with the model's lower success rate (Table IV.2). Finally, model understanding was lowest for model M_{4E}^{SQ}, slightly higher for model M_{6E}^{SQ}, and the highest for the baseline model M_7. A Kruskal-Wallis test indicated significant effects for model

Table IV.7: Descriptive Statistics by Model and Post-Trial Questionnaire Metrics.

Model	Statistic	Investigate	Abandon	Decide	Performance	Understanding
M_{4E}	Mean (SD)	5.07 (1.25)	6.04 (1.35)	5.71 (1.94)	5.11 (0.99)	5.39 (1.31)
	Median	5	6	7	5	5.5
	Min (Max)	2 (7)	1 (7)	1 (7)	3 (7)	2 (7)
M_{6E}	Mean (SD)	4.75 (1.53)	6.18 (1.42)	5.57 (1.99)	5.54 (1.29)	5.82 (1.16)
	Median	5	7	6	6	6
	Min (Max)	2 (7)	1 (7)	1 (7)	3 (7)	3 (7)
M_7	Mean (SD)	5.18 (1.68)	5.29 (1.76)	6.54 (0.92)	5.75 (1.43)	5.93 (1.46)
	Median	5.5	5.5	7	6	7
	Min (Max)	1 (7)	1 (7)	4 (7)	2 (7)	3 (7)

understanding, $\chi^2(2, N = 84) = 6.7985$, $p = 0.0334$, but did not indicate significance for the remaining metrics. Pairwise comparisons using a Tukey and Kramer (Nemenyi) test indicated significant effects for trial understanding between model M_{4E}^{SQ} and model M_7 ($p = 0.034$).

The post-experiment questionnaire required participants to rank order the models, on a scale from one (best) to three (worst). The rankings were mutually exclusive. The descriptive statistics for the post-experiment rankings are presented in Table IV.8. Lower values indicate that a model was perceived to respond more quickly to participant requests (Responsiveness), enabled more accurate target selections (Performance), or were easier for the participant to understand (Comprehension). The participants' subjective comparisons between the M_{4E}^{SQ} and M_{6E}^{SQ} were virtually equivalent across the metrics. The largest noticeable difference between the models was that the participants consistently indicated that M_7 was the easiest model to understand, but was the least responsive to their requests. While M_7 was expected to be easier to understand, since the participants controlled all aspect's of the model's decision making process, the model's perceived responsiveness was lower than expected. A Kruskal-Wallis test did not identify any significant effects for the post-experiment rankings.

Table IV.8: Descriptive Statistics by Model and Post-Experiment Questionnaire Metrics.

Model	Statistic	Responsiveness	Performance	Comprehension
M_{4E}^{SQ}	Mean (SD)	2.04 (0.84)	2.11 (0.83)	2.18 (0.82)
	Median	2	2	2
	Min (Max)	1 (3)	1 (3)	1 (3)
M_{6E}^{SQ}	Mean (SD)	1.82 (0.61)	2.11 (0.74)	2.07 (0.72)
	Median	2	2	2
	Min (Max)	1 (3)	1 (3)	1 (3)
M_7	Mean (SD)	2.11 (0.96)	1.71 (0.85)	1.68 (0.86)
	Median	2.5	1	1
	Min (Max)	1 (3)	1 (3)	1 (3)

Finally, the results of the Mental Rotation Tests (MRT-A), were compared to the collective success rates (see Chapter IV.2.2) in order to determine the existence of a correlation between spatial abilities and performance with collective action selection. Participant's MRT-A scores were compared to success rates for each model and each problem difficulty (easy and hard). A correlation was only observed in two pairs of variables. There was a positive correlation between MRT-A scores and overall success with M_{4E}^{SQ}, $r = 0.453$, $p < 0.05$, and a stronger correlation between MRT-A scores and M_{4E}^{SQ}'s success for the hard decision problems, $r = 0.532$, $p < 0.01$. These findings suggest that participants with higher spatial awareness were more capable of making accurate target selections with the M_{4E}^{SQ} model, than those participants with lower scores. The absence of a correlation between the MRT-A scores and the M_{6E}^{SQ} model suggests that model M_{6E}^{SQ} required less assistance from the human operator in making an accurate target selection decision.

IV.2.6 Discussion

The results of the human-collective site selection experiment were mixed. While the objective performance measures indicated specific advantages across the evaluated models, these advantages were less visible in the subjective measures overall. A clear example of the deviation between the objective and subjective measures was the nearly identical perceived performance between the M_{4E}^{SQ} and M_{6E}^{SQ} models in the post-trial questionnaires, despite the former model's significantly faster decision times and the latter model's higher success rates.

The results supported hypothesis $H_{5.1}$, which predicted that the presented behavior selection control mechanisms enable humans to exert a variety of levels of control over the collective action selection process. The relatively high success rates achieved by all the models supports the hypothesis and is clearest in the success rates for the M_{4E}^{SQ} model. Although this model achieved lower success rates than models M_{6E}^{SQ} and M_7, M_{4E}^{SQ}'s performance during difficult decisions was quite a bit higher than in similar decisions in the previous static, dynamic, and sequential site selection experiments (see Chapter III.3). One likely explanation for the M_{4E}^{SQ} model's higher success rate is the intervention rate, which indicated that the operator was often, although not always, able to prevent the collective from choosing a suboptimal target. This observation is further confirmed by the increased performance in both success rate and decision times when comparing the independent collective performance (C_{4E}^{SQ}) to the performance of the human-collective team (M_{4E}^{SQ}) in Fig. IV.3. Using the provided controls, the participants were clearly able to overcome the non-bias reducing model's poor independent success rate and make accurate team decisions.

Hypothesis $H_{5.2}$, which stated that the bias reducing model (M_{6E}^{SQ}) was expected to lower participant workload compared to the model lacking bias reduction mechanisms (M_{4E}^{SQ}) was only partially supported. The similar request frequencies observed for models M_{4E}^{SQ} and M_{6E}^{SQ} suggest that the bias reducing model did not reduce the frequency of user input to the collectives; however, the higher intervention rate observed for M_{4E}^{SQ} indicates a difference in participants' strategies when using each model. Model M_{6E}^{SQ} required fewer interventions to achieve a higher success rate, compared to the M_{4E}^{SQ} model. Although many of the subjective results were not significantly different between the two models, participants using M_{6E}^{SQ} were slightly more accurate in their responses to the situation awareness probes. Additionally, spatial abilities, as measured with the MRT-A, were moderately correlated with M_{4E}^{SQ}'s success rate, but not with M_{6E}^{SQ}'s success rate. This observation suggests that more participant engagement when using the M_{4E}^{SQ} model, which did not include the bias reducing mechanisms, although this phenomenon requires further investigation. The strongest subjective measure that refutes hypothesis $H_{5.2}$ is the nearly equal perceived workload between the M_{4E}^{SQ} and M_{6E}^{SQ} models. One of the challenges imposed by the experimental design was the inability to categorize most subjective measures between hard and easy target selection problems. The objective measures differed between the problem difficulties, but workload estimates, situational awareness measures, and participant evaluations of each collective were not directly associated with problem difficulty. Finally, the results partially supported hypothesis $H_{5.3}$, which predicted that the collective action selection models benefit from human interaction in terms of improved success rates and reduced decision times. The success rates of the independent collectives (C_{4E}^{SQ} and C_{6E}^{SQ}) were lower than those achieved by the human-collective teams (M_{4E}^{SQ} and M_{6E}^{SQ}) for both easy and difficult problems. Decision times generally improved as a result of the human interaction; however, the decision time of the non-bias reducing model increased with human interaction, as the participants were required to slow the collective's decision making process in order to make accurate decisions.

The fourth hypothesis, H_6, predicts that the visualization techniques, consisting of the collective state estimations and target information, are sufficient to enable participant understanding and interaction. The subjective results indicated that the hypothesis was supported; however, the technique must be compared to other visualization techniques in order for its benefit to be fully evaluated (e.g., a full representation of the collective (Walker et al., 2016) or a radial display centered on the decision making hub (Crandall et al., 2017)). Participants with all models achieved a high accuracy when responding to the situational awareness probes, indicating that the participants either already knew, or quickly determined their responses through interactions with the interface. Most trials resulted in high success rates, despite the participant's divided attention between four collectives, that changed locations during the trials. Although not conclusive, this observation indicates that the presented visualization techniques were at least sufficient to enable the participants to accomplish the evaluated task.

IV.3 Summary

This chapter introduced and evaluated new control mechanisms and visualization techniques for use in future human-collective action selection scenarios. Behavior selection control mechanisms were chosen because they were easily adapted to existing states in the individual action selection model (see Chapter III.2.3.1) and readily accessible through interaction with the collectives' decision making hubs. The presented visualization technique included a simple display that presented an estimate of the collectives' decision making process, the collectives' supported options (e.g., targets), and alerts to expected collective transitions between deliberation and action execution. The results of a human-collective site selection experiment indicated that both the presented control mechanisms and visualization techniques were effective. The experiment supported further investigation of the benefits to the human operator of interacting with discrete consensus achievement strategies with and without mechanisms for reducing environmental bias. The hypotheses tested in this chapter are summarized in Table IV.9.

Table IV.9: Summary of Research Questions and Hypotheses.

Question	Hypothesis		Experiment	Supported
R_5	$H_{5.1}$	Behavior selection mechanisms enable effective human guidance and control over explicit collective action selection.	Human-Collective Site Selection	Yes
	$H_{5.2}$	Environmental bias reduction will reduce human workload during human interaction with collective action selection.	Human-Collective Site Selection	Partially
	$H_{5.3}$	Human input improves the accuracy and decision times for the collective action selection models	Human-Collective Site Selection	Partially
R_6	H_6	The presented visualization techniques is sufficient to enable human understanding and interaction with collective action selection.	Human-Collective Site Selection	Yes

The human-collective site selection experiment provided partial answers to research questions R_5 and R_6 by evaluating participants' ability to accomplish sequences of target selection decisions with multiple collectives using the presented control mechanisms and visualization techniques. Research question R_5 was intended to determine effective control mechanisms for human-collective action selection decisions. The presented behavior control mechanisms were readily applicable to the collective action selection process and permitted the participants to manipulate multiple collectives' levels of autonomy and decision process simultaneously via interactions with the collectives' decision making hubs. The fact that the participants were successful, overall, in working with the collectives to make accurate decisions suggests that the behavior selection mechanisms are effective, but answering R_5 requires further investigation in order to compare the presented control mechanisms with other existing control mechanisms (e.g., leader selection). The second research question, R_6 was also partially addressed. Currently, the only other visualization technique proposed for use in human-collective best-of-*n* decisions is a radial display centered on the decision making hub (Crandall et al., 2017). While the radial display has potential use in both implicit and explicit action selection in human-collective teams, the technique has yet to be evaluated. The visualization technique presented in this dissertation was demonstrated to enable a human to control multiple decision processes simultaneously, which strongly supports its use in future human-collective teams. A comparison study between these existing approaches and other emergent approaches is an important area for future work.

Chapter V

Conclusion

A novel decentralized action selection strategy was developed that permits a robotic collective to choose and implement its best available action, respond to human input, and compensate for negative environmental bias. The strategy addressed three common problems that impede the use of robotic collectives in dynamic environments. First, new mechanisms reduced the influence of environmental bias, which prevents existing collective decision making strategies from choosing options based solely on value, despite differences in options' locations. Second, while existing collective task sequencing strategies have transitioned between a series of known actions, the new strategy combines this approach with collective best-of-*n* decisions to facilitate sequential action selection and execution. Finally, methods of human interaction with collectives capable of independent decisions are emergent, but untested. This dissertation introduced and evaluated new control mechanisms and visualization techniques that facilitated human-collective interaction for multiple collectives making simultaneous best-of-*n* decisions.

The developed action selection strategy consisted of a novel collective behavior model that guided the development of a variety of individual behavior models. The strategy restricted individual robot interactions to a decision making hub, which has been shown to improve collective information pooling and decision accuracy in biologically-inspired collective decision making strategies. The individual behavior models were distinguished by the use of two environmental bias reducing mechanisms, two quorum detection mechanisms, and three quorum noise reduction mechanisms.

A series of site selection experiments evaluated the effectiveness of the different mechanisms in achieving the desired emergent behaviors described by the collective model. Collectives of individuals using the bias reduction mechanisms made slower decisions, when compared to existing collective best-of-*n* models, but were better able to compensate for negative environmental bias. When required to choose options based on value, the bias reduction models compensated for option distance to the decision making hub and chose the highest valued option despite its location, while existing strategies frequently chose the closer option regardless of its value. Both the existing models and the bias reducing models were capable of breaking symmetry, meaning that each made random decisions when faced with equally valued options. The existing strategies only face symmetry breaking scenarios between equally valued, equidistant options. The bias reduction model; however, causes symmetry breaking to occur whenever two equally valued options are anywhere within the collective's search space. The quorum detection mechanisms enabled the collective to transition to action execution after making an action decision and return to a decision making state, once the action was completed. Quorum noise reduction mechanisms improved the collective's ability to complete sequences of actions without the collective making split decisions, as a result of individuals making inaccurate estimates of the collective's decision making process and prematurely initiating collective behavior transitions.

A new human-collective interaction approach permitted humans to supervise and control multiple collective action selection processes simultaneously via a simulation study. Behavior selection mechanisms were introduced that permitted humans to alter the behavior states of robots within the collectives' decision making hubs, as a means of altering collective decisions. Visualization techniques that presented estimates of the collective decision making process were shown to be sufficient for enabling humans to maintain situational awareness over multiple collectives during the experiment. The provided behavior selection control mechanisms allowed the humans to interact with the collectives' decision making processes by directing agents to support options, abandon options, or commit to executing an option. Humans interacted with action selection strategies that either reduced environmental bias, or did not. While the human-collective teams were significantly more accurate with the bias reducing strategy, the participants perceived little difference in workload or responsiveness between the strategies.

This dissertation developed two important capabilities that previously did not exist for robotic collective decision making. First, the demonstrated bias reducing mechanisms increase collective decision making resilience to undesirable environmental features. This capability will enable future robotic collectives to control the relative importance of option cost and value during decision making. Second, the presented action selection model and human-interaction methods permit the collective to participate in decisions regarding the collective's future behavior, which previously was limited to the direction of a human with assumed global awareness. The combination of these capabilities provides humans with the ability to influence a collective's consideration of the option value and cost during decision making and delegate decision making functions to the collective. These capabilities are applicable to environmental monitoring and space exploration, which will require collectives to cover wide areas, be robust to individual agent failure, prioritize information transmissions to human supervisors, and to make independent autonomous collective behavior decisions.

V.1 Contributions

Several novel contributions are presented for the fields of robotic collective decision making and Human-Swarm/Collective Interaction. The primary contributions are the explicit collective action selection strategy, an interaction delay mechanism, and the human-collective control and visualization techniques. The secondary contributions include an interaction frequency modulation mechanism and a novel quorum detection technique that, with the use of quorum noise reduction, increased the speed of the collective action selection process.

V.1.1 Primary Contributions

The first primary contribution is a novel explicit collective action selection strategy that combines collective best-of-*n* decision making with task sequencing to facilitate sequential collective behavior decisions. Existing collective decision strategies either stop after a single decision is detected (e.g., Parker and Zhang, 2009), or execute a known sequence of tasks (e.g., Parker and Zhang, 2010; Wessnitzer and Melhuish, 2003). The developed approach permits the collective to decide which of the available set of options in the environment to act on between executed tasks. A value-sensitive collective best-of-*n* decision strategy was extended with a quorum detection mechanism that retains the collective's decision making behavior, but enables the collective to detect when a decision has been reached (Cody and Adams, 2017a). A series of quorum detecting states further extends the model in order for the collective to transition successively through action initiation, execution, and back to a decision making state (Cody and Adams, 2018b). Enabling a collective to make decisions about future behaviors provides a level of autonomy that was not previously accessible to the field of Human-Swarm/Collective Interaction, which usually treats the human as the sole decision maker in a human-swarm team. Shared decision-making has been proposed (e.g., Crandall et al., 2017), but explicit action selection allows the collective to perform assessment activities and inform human decision making about future actions.

The second primary contribution is the interaction delay mechanism, which reduces negative environmental bias that was identified as an important challenge for robotic collective decision making. Collectives are often limited to local communication and sensing capabilities, a desirable feature that permits collective behavior algorithms to scale to large populations and environmental areas. Unfortunately, local interaction limitations require robots to execute movements in order to pool information and evaluate available options. Decisions that compare environmental options located at different distances from the collective's decision making hub result in higher interaction rates that bias decisions towards closer options. Interaction delays, which required robots to delay between interactive periods based on the option's distance to the hub, enabled existing collective best-of-*n* strategies to make decisions based on option value, despite environmental bias, that were not previously possible. The delays imposed longer decision times, that were mitigated by interaction frequency modulation (Cody and Adams, 2017b). Experiments also identified that negative environmental bias strongly influences the ability of existing collective best-of-*n* strategies to respond to dynamic changes in option value (Cody and Adams, 2018c). Without interaction delays, collectives responded quickly to value changes that occurred with options close to the decision making hub. Value changes to options that were further from the decision making hub; however, required much more time to reach make decisions, which were often incorrect. The developed interaction delay mechanism enabled the collective to respond uniformly with high accuracy and consistent decision time in the same scenarios.

The final primary contribution is the set of collective control and visualization techniques developed for supporting human-collective action selection decisions, which are the first techniques presented for this purpose. One challenge of interacting with collective decision strategies is that the decisions are driven by information that is not accessible (i.e., visible) to the human. The developed visualization technique provided the human with an estimate of the state of the collective decision making process using a combination of agent state summaries and informative target site symbols. This technique allowed the human participants to effectively and successfully interact with multiple simultaneous collective action selection processes. The developed behavior selection control mechanisms provided the human with the ability to direct robots to successfully change states within the collective action selection process. The *Investigate* command allowed the human to increase the evaluations of a chosen option, which usually increased the collective's support for the option. The *Abandon* command resulted in the opposite effect and reduced option support. Finally, the human used the *Decide* command in order to trigger a faster decision and accelerate the execution of a specific action. Through a combination of these behavior selection controls and the described visualization technique, the human participants were able to successfully guide and control the decision making process (Cody and Adams, 2018a).

V.1.2 Secondary Contributions

The first secondary contribution was the interaction frequency modulation mechanism. This mechanism, which increased robot interaction attempts with option distance to the decision making hub, was shown to enable faster decisions without significantly altering accuracy. Collective decision making strategies typically assume that the probability of any two agents interacting is approximately equal, which is often not the case in real-world systems (Reina et al., 2016). Collective best-of-n strategies, as previously described, often favor closer sites with higher interaction rates. While the mechanism's original purpose to counteract this bias was not fully realized in this dissertation, the mechanism significantly increased interaction rates, mitigated long decision times caused by interaction delays (see Chapter V.1.1), and increased the speed of the original collective best-of-n model (Cody and Adams, 2017b). These results suggest similar modification to interaction rates in future collective best-of-n decision making strategies.

The last secondary contribution was the episodic queue mechanism for quorum detection. Persistent queues, in which agents retained a queue of a specified length regardless of their internal state, have been evaluated in similar decision strategies (Parker and Zhang, 2009). Within the developed collective action selection strategy, these queues resulted in long, but accurate collective action selections. Episodic queues were cleared each time the agents changed state and were replenished more frequently than typical agent interactions. The episodic queues resulted in rapid decision times, comparable decision accuracy, though less accurate quorum detections, when compared to the persistent queue mechanisms (Cody and Adams, 2017a). When combined with quorum noise reduction, which reduces duplicate messages during quorum estimation, the episodic queues enabled faster sequences of collective action selections than the persistent queuing technique (Cody and Adams, 2018b).

V.2 Future Work

The developed collective action selection strategy and human-collective interaction techniques offer several areas of future work including greater exploration of the models' large parameter space, greater theoretical analysis, extensions to the presented models, and additional control and visualization techniques supporting future human-collective interaction.

V.2.1 Parameter Space Exploration

A weakness to the experimental results presented in Chapters III.3 and IV.2 are that the models were only examined within a small portion of their large available parameter space. These parameters include those associated with robot interactions, transitions between states, voting message frequency, voting queue length, communication range, and agent movement. Robot interaction rates were fixed to maintain an approximately 10% interactive population; however, exploration of time-changing latency transition periods are expected to improve decision speed. Additionally, the use of time changing inhibition rates has been proposed as a possible mechanism used by honeybees to alter decision times (Pais et al., 2013). The decision accuracies achieved by the episodic queue models, despite the models' lower quorum accuracies (see Chapter III.3.3), suggest that it is possible to shorten the collective best-of-*n* decision making strategy by lowering the quorum threshold below the evaluated value ($T = 0.75$), without reducing decision accuracy. Communication range and movement parameters affect the distribution of opinions within the search space, which causes the collective behavior to deviate significantly from the predictive model (Reina et al., 2016). Agent movement and sensing models also determine expected site discovery times, which are known to affect decision accuracy when sites are not equidistant to the nest area (Reina et al., 2015b; Cody and Adams, 2017b).

The interaction delay mechanism makes strong assumptions about the probability of site discovery within the environment. Round trip travel times and discovery rates were approximated assuming constant, homogeneous movement throughout the collective's area. Numerous factors, including robot malfunction, robot congestion, terrain slope, current or wind changes, and obstacles will impose heterogeneous movements within future environments and have not been thoroughly evaluated in opinion-based site selection scenarios. These factors have been considered in ad-hoc studies (e.g., Schmickl and Crailsheim, 2008) and in a human-collective interaction proposal (Crandall et al., 2017). One solution to these problems is to increase the collective's interaction delays in order to allow for unexpected difficulties in discovering or traveling to sites within the search space. Such an approach must consider the likelihood that travel times to specific sites will be generally consistent during the decision making scenarios, which determines the effectiveness of the delays. Additionally, increasing interaction delays more than necessary will reverse the environmental bias and encourage the selection of sites with higher costs, which is clearly undesirable.

V.2.2 Further Theoretical Analysis

Further theoretical analysis of the presented models is required to provide performance guarantees for a variety of the previously discussed parameter values and decision scenarios. Both the value-sensitive decision strategy and collective action selection are areas of continuing research. Analysis of the decision strategy has revealed the critical role of inhibition in breaking deadlock between two equally valued sites (Seeley et al., 2012) and in controlling the minimum distinguishable difference between site values (Pais et al., 2013). The limitations of the collective decision making strategy, in terms of its number of sites, number of agents, decision making hub size, and collective search area remain open areas of research. Reina *et al.* (2017) recently extended the value-sensitive decision model in order to analyze problems with three, or more sites. The work revealed that the ratio of interactive transitions (e.g., recruiting and inhibition) to spontaneous transitions (e.g., abandonment and discovery) is critical to maintaining the desired collective behavior in more complex problems (Reina et al., 2017).

V.2.3 Model Extensions

Several directions for future work exist for the collective and individual behavior models. The quorum detection mechanism for problems in which the collective must alter decision making hub locations was inadequate for reliably sensing the end of the completion of the collective's movement. The effectiveness of the sequential action selections (see Chapters III.3.4.2-III.3.4.5) largely resulted from slowing down the quorum detection mechanisms in order to prevent agents from being left behind and causing the collective to split. Solutions to this problem include the use of decentralized population or density estimates. Another area of future work is to combine a collective conducting action selection with collectives performing other tasks (e.g., foraging) and use one collective to manage behavior transitions in the other in a manner similar to honeybee scouts during nest site selection (Seeley, 2010). A final future work extension is to associate more than one action to each available option, which will enable more robust responses in complex environments.

V.2.4 Human-Collective Interaction Extensions

While the human-collective site selection experiment demonstrated the potential of future human-collective action selection, additional control and visualization techniques require further investigation. First, control mechanisms were not developed to enable the human operator to alter the search space for the collectives, although this has been proposed (Crandall et al., 2017). Such an addition permits greater use of the human's reasoning in the decision process by enabling the system to respond to other sources of information (e.g., intelligence reports and additional imagery). Second, the interface was assumed to have reliable communication with each of the decision making hubs. A more detailed simulation of the communication between the human and decision making hubs, as well as between the robots themselves, enables further investigation of the effects of communication latency and congestion, which were not examined in this dissertation. Third, mechanisms that enable the human to increase or decrease the importance of option values in relation to option locations are needed to make the system more practical. Finally, the use of heterogeneous collective populations, as discussed in Chapter II.3.4, has yet to be explored in human-collective decision strategies.

Bibliography

Abaid, N., Butail, S., Porfiri, M., and Spinello, D. (2015). Dynamics of animal systems. *The European Physical Journal Special Topics*, 224(17):3109–3117.

Adams, J. A., Goodrich, M., and Chen, J. (2018). Swarm transparency. In *ACM/IEEE International Conference on Human-robot Interaction*. In press.

Alboul, L., Saez-Pons, J., and Penders, J. (2008). Mixed human-robot team navigation in the GUARDIANS project. In *IEEE International Workshop on Safety, Security and Rescue Robotics*, pages 95–101.

Aoki, I. (1982). A simulation study on the schooling mechanism in fish. *Bulletin of the Japanese Society of Scientific Fisheries*, 48(8):1081–1088.

Arvin, F., Turgut, A. E., Bazyari, F., Arikan, K. B., Bellotto, N., and Yue, S. (2014). Cue-based aggregation with a mobile robot swarm: A novel fuzzy-based method. *Adaptive Behavior*, 22(3):189–206.

Attanasi, A., Cavagna, A., Castello, L. D., Giardina, I., Grigera, T. S., Jelic, A., Melillo, S., Parisi, L., Pohl, O., Shen, E., and Viale, M. (2014). Superfluid transport of information in turning flocks of starlings. *Nature Physics*, 10(9):691–696.

Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., and Zdravkovic, V. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237.

Barron, A. B., Gurney, K. N., Meah, L. F. S., Vasilaki, E., and Marshall, J. A. R. (2015). Decision-making and action selection in insects: Inspiration from vertebrate-based theories. *Frontiers in Behavioral Neuroscience*, 9:2–16.

Berman, S., Halász, A., Kumar, V., and Pratt, S. (2007). Algorithms for the analysis and synthesis of a bio-inspired swarm robotic system. In *International Conference on Swarm Robotics*, pages 56–70. Springer.

Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, Inc., New York, NY, USA.

Bonabeau, E., Theraulaz, G., and Deneubourg, J.-L. (1998). Fixed response thresholds and the regulation of division of labor in insect societies. *Bulletin of Mathematical Biology*, 60(4):753–807.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41.

Brown, D. S., Goodrich, M. A., Jung, S.-Y., and Kerman, S. (2016). Two invariants of human-swarm interaction. *Journal of Human Robot Interaction*, 5(1):1–31.

Campbell, M., D'Andrea, R., Schneider, D., Chaudhry, A., Waydo, S., Sullivan, J., Veverka, J., and Klochko, A. (2003). RoboFlag games using systems based, hierarchical control. In *American Control Conference*, volume 1, pages 661–666.

Campo, A., Garnier, S., Dédriche, O., Zekkri, M., and Dorigo, M. (2011). Self-organized discrimination of resources. *PLOS ONE*, 6(5):1–7.

Campo, A., Gutiérrez, Á., Nouyan, S., Pinciroli, C., Longchamp, V., Garnier, S., and Dorigo, M. (2010). Artificial pheromone for path selection by a foraging swarm of robots. *Biological Cybernetics*, 103(5):339–352.

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

Chamberlain, L., Tang, J., Watugala, M., Adams, J. A., and Babish, M. (2003). A behavioral architecture for strategy execution in the RoboFlag game. In *American Control Conference*, pages 672–677.

Chen, J. Y. C., Barnes, M. J., Selkowitz, A. R., and Stowers, K. (2016). Effects of agent transparency on human-autonomy teaming effectiveness. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 1838–1843.

Chen, J. Y. C., Barnes, M. J., Wright, J. L., Stowers, K., and Lakhmani, S. G. (2017). Situation awareness-based agent transparency for human-autonomy teaming effectiveness. *Proc.SPIE*, 10194:1 – 6.

Codling, E. A., Plank, M. J., and Benhamou, S. (2008). Random walk models in biology. *Journal of The Royal Society Interface*, 5(25):813–834.

Cody, J. R. and Adams, J. A. (2017a). An evaluation of quorum sensing mechanisms in collective value-sensitive site selection. In *International Symposium on Multi-Robot and Multi-Agent Systems*, pages 40–47.

Cody, J. R. and Adams, J. A. (2017b). Reducing distance bias in value-sensitive swarm site selection. In *Proceedings of the International Symposium on Swarm Behavior and Bio-Inspired Robots*, pages 223–230.

Cody, J. R. and Adams, J. A. (2018a). Human-collective action selection decisions. *In Preparation*.

Cody, J. R. and Adams, J. A. (2018b). Sequential collective action selection in value-sensitive site selection problems. *In Preparation*.

Cody, J. R. and Adams, J. A. (2018c). Value-sensitive collective action selection with dynamic option values. *Manuscript Submitted*.

Conradt, L. and Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology and Evolution*, 20(8):449 – 456.

Conradt, L. and Roper, T. J. (2009). Conflicts of interest and the evolution of decision sharing. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1518):807–819.

Coppin, G. and Legras, F. (2012). Autonomy spectrum and performance perception issues in swarm supervisory control. *Proceedings of the IEEE*, 100(3):590–603.

Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1):36 – 43.

Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(1):513–516.

Couzin, I. D., Krause, J., James, R., Graeme, D. R., and Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1):1 – 11.

Crandall, J. W., Anderson, N., Ashcraft, C., Grosh, J., Henderson, J., McClellan, J., Neupane, A., and Goodrich, M. A. (2017). Human-swarm interaction as shared control: Achieving flexible fault-tolerant systems. In Harris, D., editor, *Engineering Psychology and Cognitive Ergonomics: Performance, Emotion and Situation Awareness*, pages 266–284. Springer International Publishing.

Dorigo, M., Floreano, D., Gambardella, L. M., Mondada, F., Nolfi, S., Baaboura, T., Birattari, M., Bonani, M., Brambilla, M., Brutschy, A., Burnier, D., Campo, A., Christensen, A. L., Decugniere, A., Caro, G. D., Ducatelle, F., Ferrante, E., Forster, A., Gonzales, J. M., Guzzi, J., Longchamp, V., Magnenat, S., Mathews, N., de Oca, M. M., O'Grady, R., Pinciroli, C., Pini, G., Retornaz, P., Roberts, J., Sperati, V., Stirling, T., Stranieri, A., Stutzle, T., Trianni, V., Tuci, E., Turgut, A. E., and Vaussard, F. (2013). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics Automation Magazine*, 20(4):60–71.

Ducatelle, F., Di Caro, G. A., Pinciroli, C., and Gambardella, L. M. (2011). Self-organized cooperation between robotic swarms. *Swarm Intelligence*, 5(2):73.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64.

Endsley, M. R. and Kaber, D. B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics*, 42(3):462–492.

Ferrante, E., Turgut, A. E., Huepe, C., Stranieri, A., Pinciroli, C., and Dorigo, M. (2012). Self-organized flocking with a mobile robot swarm: A novel motion control method. *Adaptive Behavior*, 20(6):460–477.

Franks, N. R., Pratt, S. C., Mallon, E. B., Britton, N. F., and Sumpter, D. J. (2002). Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 357(1427):1567–1583.

Garnier, S., Gautrais, J., Asadpour, M., Jost, C., and Theraulaz, G. (2009). Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2):109–133.

Golman, R., Hagmann, D., and Miller, J. H. (2015). Polya's bees: A model of decentralized decision-making. *Science Advances*, 1(8):1–4.

Goodrich, M., Kerman, S., and Jun, S.-Y. (2012). *On Leadership and Influence in Human-Swarm Interaction*, pages 32–37. Technical report FS-12-04 (Association for the Advancement of Artificial Intelligence). AAAI Press.

Goodrich, M. A., Sujit, P. B., Kerman, S., Pendleton, B., and Pinto, J. (2011). Enabling human interaction with bio-inspired robot teams: Topologies, leaders, predators, and stakeholders. Technical Report 2011-1, Brigham Young University, Human Centered Machine Intelligence Laboratory.

Hamann, H., Schmickl, T., Wrn, H., and Crailsheim, K. (2012). Analysis of emergent symmetry breaking in collective decision making. *Neural Computing and Applications*, 21(2):207–218.

Hamar, J. and Dove, R. (2012). A quorum sensing pattern for multi-agent self-organizing security systems. In *IEEE International Carnahan Conference on Security Technology*, pages 319–324.

Haque, M., Ren, C., Baker, E., Kirkpatrick, D., and Adams, J. A. (2016). Analysis of swarm communication models. In *International Workshop on Combinations of Intelligent Methods and Applications*, pages 29–37.

Hayes, S. T. and Adams, J. A. (2014). Human-swarm interaction: Sources of uncertainty. In *ACM/IEEE International Conference on Human-robot Interaction*, pages 170–171, New York, NY, USA. ACM.

Humphrey, C. M. (2009). Information abstraction visualization for human-robot interaction.

Jung, S. Y., Brown, D. S., and Goodrich, M. A. (2013). Shaping Couzin-like torus swarms through coordinated mediation. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 1834–1839.

Jung, S. Y. and Goodrich, M. A. (2013). Multi-robot perimeter-shaping through mediator-based swarm control. In *International Conference on Advanced Robotics*, pages 1–6.

Kaadan, A., Refai, H. H., and LoPresti, P. G. (2014). Multielement FSO transceivers alignment for inter-UAV communications. *Journal of Lightwave Technology*, 32(24):4785–4795.

Kampen, N. G. V. (2007). Chapter II Random Events. In Kampen, N. V., editor, *Stochastic Processes in Physics and Chemistry (Third Edition)*, North-Holland Personal Library, page 171. Elsevier, third edition.

Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., and Schmickl, T. (2015). Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. In Chen, Q., Torroni, P., Villata, S., Hsu, J., and Omicini, A., editors, *Principles and Practice of Multi-Agent Systems*, pages 201–217, Cham. Springer International Publishing.

Kira, Z. and Potter, M. A. (2009). Exerting human control over decentralized robot swarms. In *International Conference on Autonomous Robots*, pages 566–571.

Kolling, A., Sycara, K., Nunnally, S., and Lewis, M. (2013). Human-swarm interaction: An exerimental study of two types of interaction with foraging swarms. *Journal of Human Robot Interaction*, 2(2):103–129.

Kolling, A., Walker, P., Chakraborty, N., Sycara, K., and Lewis, M. (2016). Human interaction with robot swarms: A survey. *IEEE Transactions on Human-Machine Systems*, 46(1):9–26.

Kouvaros, P. and Lomuscio, A. (2016). Formal verification of opinion formation in swarms. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1200–1208.

Krieger, M. J. B. and Billeter, J.-B. (2000). The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30(1):65 – 84.

Majumdar, A. K. (2015). *Free-Space Optical (FSO) Platforms: Unmanned Aerial Vehicle (UAV) and Mobile*, pages 203–225. Springer New York, New York, NY.

Mallon, E., Pratt, S., and Franks, N. (2001). Individual and collective decision-making during nest site selection by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 50(4):352–359.

Marketer, D. B. (2017). 13 best drones with the longest flight time-2018 buyer's guide.

Mathews, N., Christensen, A. L., O'Grady, R., and Dorigo, M. (2012). Spatially targeted communication and self-assembly. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2678–2679.

Mermoud, G., Matthey, L., Evans, W. C., and Martinoli, A. (2010). Aggregation-mediated collective perception and action in a group of miniature robots. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, pages 599–606, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Mermoud, G., Upadhyay, U., Evans, W. C., and Martinoli, A. (2014). Top-down vs. Bottom-up model-based methodologies for distributed control: A comparative experimental study. *The International Symposium on Experimental Robotics*, pages 615–629.

Miller, C. (2012). Delegation management versus the swarm: A matchup with two winners. In Lewis, M. and for the Advancement of Artificial Intelligence, A., editors, *Human Control of Bioinspired Swarms: Papers from the AAAI Fall Symposium*, Technical report FS-12-04 (Association for the Advancement of Artificial Intelligence), pages 44–46. AAAI Press.

Miller, C., Funk, H., Wu, P., Goldman, R., Meisner, J., and Chapman, M. (2005). The Playbook approach to adaptive automation. *Human Factors and Ergonomics Society Annual Meeting*, 49(1):15–19.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.-C., Floreano, D., and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In Gonalves, P. J., Torres, P. J., and Alves, C. M., editors, *Proceedings of the Conference on Autonomous Robot Systems and Competitions*, volume 1, pages 59–65, Castelo Branco, Portugal. IPCB: Instituto Politcnico de Castelo Branco.

Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3-4):305–327.

Nagavalli, S., Chakraborty, N., and Sycara, K. (2017). Automated sequencing of swarm behaviors for supervisory control of robotic swarms. In *IEEE International Conference on Robotics and Automation*, pages 2674–2681.

Nevai, A. L., Passino, K. M., and Srinivasan, P. (2010). Stability of choice in the honey bee nest-site selection process. *Journal of Theoretical Biology*, 263(1):93 – 107.

O'Grady, R., Pinciroli, C., Christensen, A. L., and Dorigo, M. (2009). Supervised group size regulation in a heterogeneous robotic swarm. In *Proceedings of the Conference on Mobile Robots and Competitions*, pages 113–119.

Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., and Marshall, J. A. R. (2013). A mechanism for value-sensitive decision-making. *PLOS ONE*, 8(9):1–9.

Parasuraman, R., Galster, S., and Miller, C. (2003). Human control of multiple robots in the RoboFlag simulation environment. In *IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3232–3237 vol.4.

Parasuraman, R., Sheridan, T. B., and Wickens, C. D. (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):286–297.

Parker, C. A. C. and Zhang, H. (2009). Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251.

Parker, C. A. C. and Zhang, H. (2010). Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4(3):199–220.

Parker, C. A. C. and Zhang, H. (2011). Biologically inspired collective comparisons by robotic swarms. *The International Journal of Robotics Research*, 30(5):524–535.

Passino, K. M. (2010). Honey bee swarm cognition: Decision-making performance and adaptation. *International Journal of Swarm Intelligence Research*, 1(2):80–97.

Penders, J. (2007). Robot swarming applications. In Donkers, J., editor, *Liber Amicorum ter gelegenheid van de 60e verjaardag van Prof.dr. H. Jaap van den Herik*, pages 227–234. Maastricht ICT Competence Center, Maastricht, Netherlands. Copyright: The author.

Peters, M., Laeng, B., Latham, K., Jackson, M., Zaiyouna, R., and Richardson, C. (1995). A redrawn Vandenberg and Kuse mental rotations test-different versions and factors that affect performance. *Brain and cognition*, 28(1):39–58.

Pinciroli, C., O'Grady, R., Christensen, A. L., and Dorigo, M. (2009). Self-organised recruitment in a heteregeneous swarm. In *International Conference on Advanced Robotics*, pages 1–8.

Pinciroli, C., O'Grady, R., Christensen, A. L., and Dorigo, M. (2010). Coordinating heterogeneous swarms through minimal communication among homogeneous sub-swarms. In Dorigo, M., Birattari, M., Di Caro, G. A., Doursat, R., Engelbrecht, A. P., Floreano, D., Gambardella, L. M., Gro\s s, R., ahin, E., Sayama, H., and Sttzle, T., editors, *Swarm Intelligence*, pages 558–559, Berlin, Heidelberg. Springer Berlin Heidelberg.

Pratt, S. C., Mallon, E. B., Sumpter, D. J., and Franks, N. R. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52(2):117–127.

Pratt, S. C. and Sumpter, D. J. T. (2006). A tunable algorithm for collective decision-making. *National Academy of Sciences*, 103(43):15906–15910.

Proud, R. W., Hart, J. J., and Mrozinski, R. B. (2003). Methods for determining the level of autonomy to design into a human spaceflight vehicle: A function specific approach. In *Performance Metrics for Intelligent Systems*, pages 1–15.

Reas, C. and Fry, B. (2014). *Processing: A Programming Handbook for Visual Designers and Artists*. MIT Press.

Reina, A., Bose, T., Trianni, V., and Marshall, J. (2016). Effects of spatiality on value-sensitive decisions made by robot swarms. In *International Symposium on Distributed Autonomous Robotic Systems*, pages 1–16. Springer.

Reina, A., Dorigo, M., and Trianni, V. (2014). Towards a cognitive design pattern for collective decision-making. *In: Dorigo M. et al. (eds) Swarm Intelligence. Lecture Notes in Computer Science*, 8667:194–205.

Reina, A., Marshall, J. A. R., Trianni, V., and Bose, T. (2017). Model of the best-of-*n* nest-site selection process in honeybees. *Physical Review Letters*, 95:052411.

Reina, A., Miletitch, R., Dorigo, M., and Trianni, V. (2015a). A quantitative micromacro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102.

Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., and Trianni, V. (2015b). A design pattern for decentralised decision making. *PLoS ONE*, 10(10):1–18.

Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH Computer Graphics*, 21(4):25–34.

Ross-Gillespie, A. and Kummerli, R. (2014). Collective decision-making in microbes. *Frontiers in Microbiology*, 5(54):1664–302X (Electronic).

Rubenstein, M., Ahler, C., and Nagpal, R. (2012). Kilobot: A low cost scalable robot system for collective behaviors. In *IEEE International Conference on Robotics and Automation*, pages 3293–3298.

Şahin, E. (2005). Swarm robotics: From sources of inspiration to domains of application. *In: Şahin E., Spears W.M. (eds) Swarm Robotics. Lecture Notes in Computer Science*, 3342. Springer, Berlin, Heidelberg:10–20.

Sahingoz, O. K. (2014). Networking models in Flying Ad-hoc Networks (FANETs): Concepts and challenges. *Journal of Intelligent & Robotic Systems*, 74(1):513–527.

Sasaki, T. and Pratt, S. C. (2012). Groups have a larger cognitive capacity than individuals. *Current Biology*, 22(19):827 – 829.

Sasaki, T. and Pratt, S. C. (2013). Ants learn to rely on more informative attributes during decision-making. *Biology Letters*, 9(6):1–4.

Scheidler, A., Brutschy, A., Ferrante, E., and Dorigo, M. (2016). The k-unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics*, 46(5):1175–1188.

Schmickl, T. and Crailsheim, K. (2008). Trophallaxis within a robotic swarm: Bio-inspired communication among robots in a swarm. *Autonomous Robots*, 25(1):171–188.

Seeley, D. T. and Buhrman, C. S. (2001). Nest-site selection in honey bees: How well do swarms implement the "best-of-n" decision rule? *Behavioral Ecology and Sociobiology*, 49(5):416–427.

Seeley, T., Visscher, P., Schlegl, T., Hogan, P., Franks, N., and Marshall, J. (2012). Stop signals provide cross inhibition in collective decision-making by honey bee swarms. *Science*, 335:108–111.

Seeley, T. D. (1982). How honeybees find a home. *Scientific American*, 247(4):158–168.

Seeley, T. D. (1995). *The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies*. Havard University Press, Cambridge, Massachusetts, London, England.

Seeley, T. D. (2010). *Honeybee Democracy*. The Princeton University Press, 1st edition.

Seeley, T. D., Camazine, S., and Sneyd, J. (1991). Collective decision-making in honey bees: How colonies choose among nectar sources. *Behavioral Ecology and Sociobiology*, 28(4):277–290.

Seeley, T. D. and Morse, R. A. (1978). Nest site selection by the honey bee, *Apis mellifera*. *Insectes Sociaux*, 25(4):323–337.

Seiffert, A., Hayes, S. T., Harriott, C. E., and Adams, J. A. (2015). Motion perception of biological swarms. In *Annual Meeting of the Cognitive Science Society*, pages 2128–2133.

Selcon, S., Taylor, R., and Koritsas, E. (1991). Workload or situational awareness?: TLX vs. SART for aerospace systems design evaluation. In *Proceedings of the Human Factors Society Annual Meeting*, volume 35, pages 62–66. SAGE Publications Sage CA: Los Angeles, CA.

Shaffer, Z., Sasaki, T., and Pratt, S. C. (2013). Linear recruitment leads to allocation and flexibility in collective foraging by ants. *Animal Behaviour*, 86(5):967 – 975.

Sheridan, T. B. and Verplank, W. L. (1978). Human and computer control of undersea teleoperators (man-machine systems laboratory report).

Spears, W., Spears, D., Hamann, J., and Heil, R. (2004). Distributed, physics-based control of swarms of vehicles. *Autonomous Robots*, 17(2-3):137–162.

Strandburg-Peshkin, A., Twomey, C. R., Bode, N. W., Kao, A. B., Katz, Y., Ioannou, C. C., Rosenthal, S. B., Torney, C. J., Wu, H. S., Levin, S. A., and Couzin, I. D. (2013). Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23(17):R709 – R711.

Stranieri, A., Ferrante, E., Turgut, A. E., Trianni, V., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). Self-organized flocking with an heterogeneous mobile robot swarm. In *ECAL*.

Sumpter, D. J. and Pratt, S. C. (2009). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1518):743–753.

Sumpter, D. J. T. (2010). *Collective Animal Behavior*. Princeton University Press, 1st edition.

Swain, D. T., Couzin, I. D., and Leonard, N. E. (2015). Coordinated speed oscillations in schooling killifish enrich social communication. *Journal of Nonlinear Science*, 25(5):1077–1109.

Theraulaz, G., Gautrais, J., Camazine, S., and Deneubourg, J.-L. (2003). The formation of spatial patterns in social insects: From simple behaviours to complex structures. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 361(1807):1263–1282.

Truszkowski, W., Hinchey, M., Rash, J., and Rouff, C. (2004). Nasa's swarm missions: The challenge of building autonomous software. *IT Professional*, 6(5):47–52.

Valentini, G. (2017). *Achieving Consensus in Robot Swarms: Design and Analysis of Strategies for the best-of-n Problem*. Studies in Computational Intelligence. Springer International Publishing.

Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9.

Valentini, G., Hamann, H., and Dorigo, M. (2014). Self-organized collective decision making: The weighted voter model. In *International Conference on Autonomous Agents and Multi-agent Systems*, pages 45–52.

Valentini, G., Hamann, H., and Dorigo, M. (2015). Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1305–1314.

Vandenberg, S. G. and Kuse, A. R. (1978). Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills*, 47(2):599–604.

Walker, P., Amirpour Amraii , S., Chakraborty, N., Lewis, M., and Sycara, K. (2014). Human control of robot swarms with dynamic leaders. In *International Conference on Intelligent Robots and Systems*, pages 1108–1113.

Walker, P., Lewis, M., and Sycara, K. (2016). The effect of display type on operator prediction of future swarm states. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 002521–002526.

Walker, P., Nunnally, S., Lewis, M., Chakraborty, N., and Sycara, K. (2013). Levels of automation for human influence of robot swarms. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 57(1):429–433.

Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., and Sycara, K. (2012a). Neglect Benevolence in Human Control of Swarms in the Presence of Latency. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 3009–3014.

Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., and Sycara, K. (2012b). Neglect benevolence in human control of swarms in the presence of latency. pages 3009–3014.

Ward, A. J. W., Sumpter, D. J. T., Couzin, I. D., Hart, P. J. B., and Krause, J. (2008). Quorum decision-making facilitates information transfer in fish shoals. *National Academy of Sciences*, 105(19):6948–6953.

Wessnitzer, J. and Melhuish, C. (2003). Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In Banzhaf, W., Ziegler, J., Christaller, T., Dittrich, P., and Kim, J. T., editors, *Advances in Artificial Life*, pages 893–902, Berlin, Heidelberg. Springer Berlin Heidelberg.

Wilson, E. O. (1984). The relation between caste ratios and division of labor in the ant genus *Pheidole Hymenoptera: Formicidae*. *Behavioral Ecology and Sociobiology*, 16(1):89–98.

Young, G. F., Scardovi, L., Cavagna, A., Giardina, I., and Leonard, N. E. (2013). Starling flock networks manage uncertainty in consensus at low cost. *PLoS Computational Biology*, 9(1):e1002894.