Heterogeneous Response Intensity Ranges and Response Probability Improve Goal Achievement in Multi-agent Systems

H. David Mathias[1](✉), Annie S. Wu[2], and Laik Ruetten[1]

[1] University of Wisconsin - La Crosse, La Crosse, WI, USA
dmathias@uwlax.edu
[2] University of Central Florida, Orlando, FL, USA
aswu@cs.ucf.edu

Abstract. Inter-agent variation is well-known in both the biology and computer science communities as a mechanism for improving task selection and swarm performance for multi-agent systems. Response threshold variation, the most commonly used form of inter-agent variation, desynchronizes agent actions allowing for more targeted agent activation. Recent research using a less common form of variation, termed dynamic response intensity, demonstrates that modeling levels of agent experience or varying physical attributes and using these to allow some agents to perform tasks more efficiently or vigorously, significantly improves swarm goal achievement when used in conjunction with response thresholds. Dynamic intensity values vary within a fixed range as agents activate for tasks. We extend previous work by demonstrating that adding another layer of variation to response intensity, in the form of heterogeneous ranges for response intensity values, provides significant performance improvements when response is probabilistic. Heterogeneous intensity ranges break the coupling that occurs between response thresholds and response intensities when the intensity range is homogeneous. The decoupling allows for increased diversity in agent behavior.

1 Introduction

Swarms of artificial agents, which model, among other things, natural colonies of insects, are comprised of some number of software or hardware agents working in concert to achieve a goal. The agents accomplish the goal by completing, usually repeatedly, one or more tasks. The swarms in this work are decentralized. Thus, there is no leader or central control of the swarm and the agents do not communicate. Each agent chooses which tasks to perform and when. Work modeling natural swarms with artificial swarms dates back two decades [14].

Agents determine which tasks to undertake by considering environmental stimuli. When agents respond to these stimuli in the same way, their actions are synchronized. This synchrony often results in poor goal achievement. Swarm performance can be improved via inter-agent variation: differences in how and

M. Dorigo et al. (Eds.): ANTS 2020, LNCS 12421, pp. 148–160, 2020.
https://doi.org/10.1007/978-3-030-60376-2_12

when agents select and perform tasks. Common forms of inter-agent variation include response thresholds [6,15] and response probabilities [15].

Response thresholds desynchronize agents' actions by varying the stimulus required for an agent to act. This models the non-determinism inherent in natural swarms. Systems typically utilize response thresholds in one of two ways: probabilistically or deterministically. Probabilistic response, introduced by Bonabeau, *et al.* [1,2], uses a formula based on a task stimulus τ and an agent's response threshold θ to determine whether the agent activates for the task. The probability of activation increases with τ, approaching 1.0 when $\tau \gg \theta$. When $\tau = \theta$, the probability is 0.5. Deterministic response [7,13,20] activates an agent if $\tau \geq \theta$. Agent actions are desynchronized only if threshold values are heterogeneous. In the biology literature, Weidenmüller [15] suggests that use of heterogeneous response thresholds together with probabilistic response can further improve diversity in agent behaviors. Wu, *et al.* studied this effect in artificial swarms, confirming the benefit of non-determinism with heterogeneity [19].

Isolation of probabilistic response into a separate form of inter-agent variation, one that can be set and tuned independent of response thresholds, leads to a form of inter-agent variation known as *response probability* [13,17–19]. A first-order effect of decreased response probability is a decrease in the number of agents that activate for a task. Perhaps more importantly, a second-order effect is that inaction by frequent actors, agents with low response thresholds, may allow other agents to gain experience with a task [15,19]. This results from increased need due to the reduction in agents performing the task. Increased need exceeds the response threshold for additional agents, allowing them to participate. The experience gained by these agents may be important to the swarm if, for example, an extinction eliminates frequent actors.

The need for redundancy in artificial swarms has been acknowledged for many years as a way to mitigate the effects of agent failure or loss [4]. Similar effects are common in natural swarms in which agents are lost due to age, predators, or competitors. If these agents are frequent actors for a task, their loss may create significant short-term difficulty for the swarm as less experienced agents must fill the void. If, however, frequent actors sometimes remain idle due to decreased response probability, agents with higher response thresholds would gain experience with the task, making the swarm more tolerant of agent loss.

Response intensity is a less known form of inter-agent variation, particularly for artificial swarms. Response intensity models differences in quantities such as a natural agent's physical size, strength or stamina, attributes that may allow the agent to work more vigorously or more efficiently. Biologists have documented variation in response intensity [3,11]. For example, in the natural world some insects are known to change their response intensity as necessary to meet the needs of their colony [5]. Response intensity may also model an agent's experience on a task. We are not aware of previous work, prior to this year, that attempts to model this natural phenomenon in artificial swarms [10].

Experience is known to impact not only individual task efficiency but also individual task selection as well as collective colony performance [9,12]. In

Cerapachys biroi ants, individuals that find early success in foraging activities choose to forage again, whereas those individuals that were unsuccessful are more likely to choose to care for young in the nest [12]. In *Leptothorax albipennis* ants, task repetition improved colony performance for emigration, the task of moving the colony to a new nesting location [9]. Because the entire colony is exposed during emigration, and therefore in danger, efficient emigration is highly desirable.

In artificial agents, response intensity may model a decrease in output due to wear and tear or the increase in the output of a new and improved device. Importantly, heterogenous response intensities, when paired with heterogeneous response thresholds, play a role in determining which agents undertake a task and, therefore, gain experience and proficiency in that task.

Mathias *et al.* [10] demonstrated that dynamic, heterogeneous response intensities significantly improve swarm task achievement when combined with heterogeneous response thresholds and result in increased agent specialization. Dynamic response intensities vary within a specified range over the course of a run, increasing when an agent activates for a task and decreasing when it does not, modeling an agent's experience with the task. The range within which response intensities vary is homogeneous.

One consequence of combining heterogeneous response thresholds with dynamic, heterogeneous response intensities is that the values couple. This occurs because agents with low thresholds for a task activate more frequently for that task. Each activation increases the response intensity for the task, within the specified range. Thus, over time, an agent's response threshold for a task correlates with its response intensity for that task. Further, if the work performed by frequent actors is sufficient to meet task demand, agents with higher thresholds are denied the opportunity to gain experience for that task. This is potentially harmful to the swarm.

In this work, we demonstrate that using dynamic, heterogeneous response intensities that vary within *heterogeneous ranges*, rather than homogeneous ranges, improves swarm performance as response probability decreases. This occurs because heterogeneous intensity ranges and decreased response probability serve to *decouple* response thresholds and response intensities. We show that this makes the swarm more resistant to the effects of extinctions of experienced agents.

2 Model and Testbed Problem

We extend previous work on response intensities in two significant ways. First, we augment the dynamic, heterogeneous response intensities with heterogeneous intensity ranges. Thus, rather than all agents sharing a common range within which their intensities vary with experience, each agent has a unique intensity range. Second, we incorporate response probability. Response probability allows an agent to fail to undertake a task when the agent's response threshold for that task is met. The response probability values used here are homogeneous.

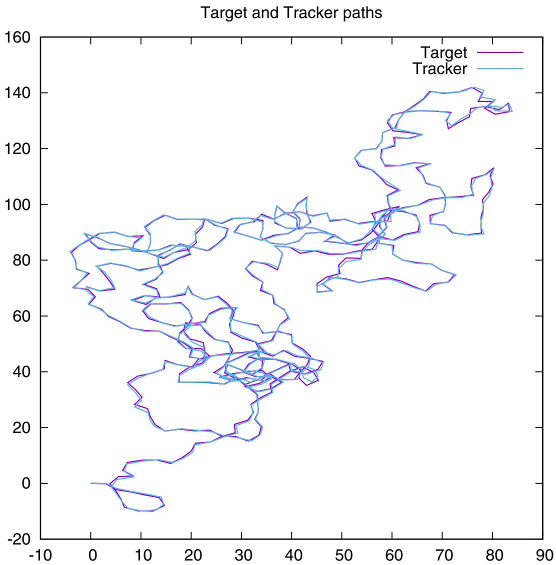

Fig. 1. An example random target path (purple) and corresponding tracker movement (blue) over 500 time steps. (Colour figure online)

Our testbed is a 2D tracking problem. This problem consists of a *target* object that moves in the plane and a *tracker* object. A swarm controls the tracker, pushing it to stay as close as possible to the target, which moves at random or according to one of several predefined paths. The paths are unknown to the agents. Each agent is capable of performing all tasks required of the swarm. The tasks are: `push_N`, `push_E`, `push_S`, or `push_W`. Agents may also remain idle if none of their response thresholds are met or due to the response probability. An example random target path is illustrated in Fig. 1.

A simulation consists of a predefined number of time steps. The target moves a fixed distance in each time step. The target's direction of travel can change as often as each time step, allowing frequently changing task demands. Agents are aware of the distance from the tracker to the target, defined by: $\Delta x = \text{target}.x - \text{tracker}.x$ and $\Delta y = \text{target}.y - \text{tracker}.y$. In each time step, each agent chooses a task to perform from among those tasks for which the agent's response thresholds are met.

Swarm goal achievement is measured according to two criteria:

Goal 1. Minimize the average positional difference, per time step, between the target location and the tracker location.

Goal 2. Minimize the difference between total distance traveled by target and the total distance traveled by the tracker.

We note that neither criterion alone is sufficient to gauge the swarm's success. Consider using only Goal 1. The tracker could remain close to the target while alternately racing ahead or falling behind. This would result in a good average difference but a path length that is significantly greater than that traveled by the target. Alternately, the tracker could travel a path that is the same length as that of the tracker while straying quite far, taking a very different path.

Swarm efficiency is measured by the number of times agents switch from one task to another and the number of agents that activate for a task in a time step. Both task switches and activations may have costs in real-world applications, thus, a swarm is more efficient when these quantities are reduced. For example, undertaking a new task might require an agent to move to a new location, incurring costs in time and fuel.

Here we define the forms of inter-agent variation used in this work. Let $a_i,\ i \in \{1, \ldots, n\}$ be an agent.

- **Response threshold**: A value $\theta_{i,D}$ $(D \in \{N, E, S, W\})$ for each task that represents the maximum acceptable Δ_D between the target and tracker for that task. If Δ_D exceeds $\theta_{i,D}$, agent a_i may activate for that `push_D`. These values are heterogeneous. Response thresholds are assigned uniformly at random in [0.0..1.0], a choice supported in the literature [7,8,16].
- **Response intensity**: A multiplier $\gamma_{i,D}$ for each task. It represents the factor by which the experience of a_i for task `push_D` differs from the default value of 1.0. This manifests as increased/decreased pushing power, equal to $\gamma_{i,D}$. These values are dynamic and heterogeneous. They are initialized uniformly at random within the agent's response intensity range for that task.
- **Response intensity range**: Intensity multipliers increase or decrease with an agent's experience for a task. The values for task `push_D` for agent a_i are bounded within a range $[\gamma_{i,D}min, \gamma_{i,D}max]$. Ranges may be homogeneous or heterogeneous. See Sect. 3 for a more detailed discussion.
- **Response probability**: A value p that represents the probability that an agent activates for a task. This value is homogeneous across all agents and tasks. It is a parameter to our system and is varied between runs.

3 Experimental Design

As response intensity and response probability are the focus of this work, we run experiments with two different types of intensity ranges – homogeneous and heterogeneous – and 7 response probability values, [0.4..1.0] in increments of 0.1.

To model the loss of agents and our system's ability to recover from such events, we implement three different forms of agent extinction. `kill-0`, in which no agents are killed; `kill-20-100-0`, in which 20 agents are killed at time step 100; and `kill-20-100-100`, in which 20 agents are killed every 100 time steps beginning at time step 100. In each case, the agents chosen for extinction are those that were idle in the fewest time steps. This means that we kill those agents with the most experience and examine how well the swarm is able to recover.

Homogeneous intensity ranges are fixed at $[0.5, 2.0]$. Heterogeneous intensity range for agent a_i and task `push_D` is assigned by first choosing a size d uniformly

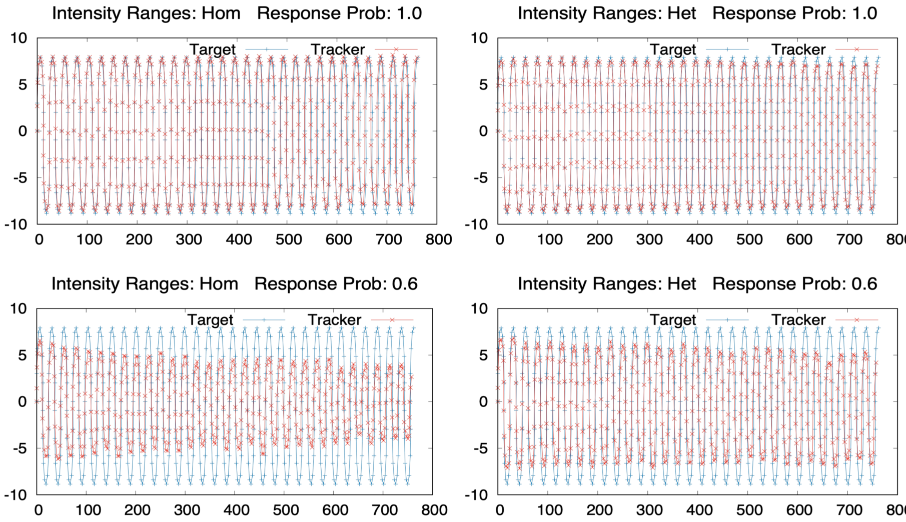

Fig. 2. A tracker performance comparison for four representative runs on path s-curve. The rows differ in response probability with 1.0 above 0.6. The columns are different response intensity ranges with homogeneous followed by heterogeneous. With response probability 0.6, the tracker performs substantially better with heterogeneous ranges. (Color figure online)

at random in $[0.6, 1.6]$. Offset $\gamma_{i,D}min$ is then chosen uniformly at random in $[0.3, (2.4 - d)]$. $\gamma_{i,D}max = \gamma_{i,D}min + d$ yielding range $[\gamma_{i,D}min, \gamma_{i,D}max] \subset [0.3, 2.4]$. These values are also determined empirically under the same conditions listed above. We note that the upper and lower endpoint values for both homogeneous and heterogeneous ranges are empirically determined to optimize behavior for the respective intensity range type for runs in which the response probability is 1.0 and no agent extinctions occur.

We test our system on three target paths: random, s-curve, and sharp. Random paths are generated by calculating an angle change, in radians, at every time step. The change is Gaussian $\mathcal{N}(0.0, 1.0)$. S-curve is a periodic curve seen in Fig. 2. Sharp is a randomized path in which a new heading and probability q of changing direction are chosen in every time step. The heading is chosen uniformly in $[0, 2\pi]$ and q is uniform in $[0.2, 0.6]$. Thus, turns are sharper than in the random path. All three paths create changing task demands though, random and sharp change more dramatically.

The variations discussed in this section produce 126 experiments for testing, 42 for each of the target paths. For each experiment we perform 100 runs. Each run lasts 500 time steps. In each time step, the target moves 3 distance units for a total path length of 1500. The swarm consists of 200 agents each of which is capable of performing all four tasks.

At each time step, we record the tracker's distance from the target. In addition, we record the total distance traveled by the target, total distance traveled by the tracker, the number of time steps in which each agent pushes in each direction, the number of times an agent does not perform a task (remains idle), and the number of times an agent switches from one task to another.

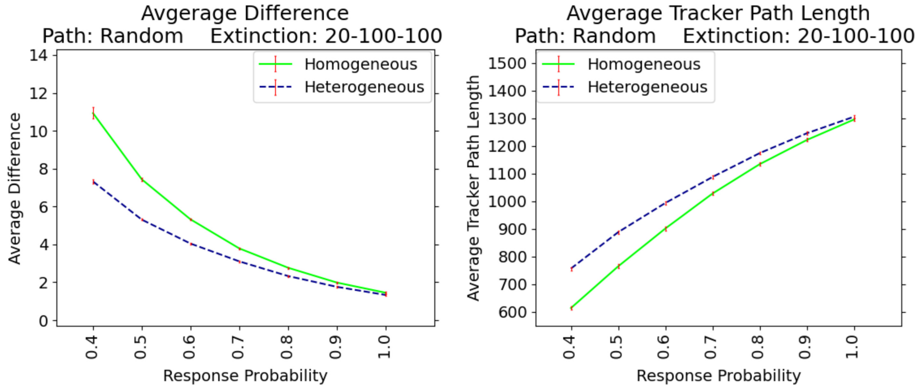

Fig. 3. Average positional difference and tracker path length for homogeneous and heterogeneous intensity ranges for target path random for 100 runs. Error bars are shown in red. Both quantities are improved for heterogeneous ranges. (Color figure online)

4 Experimental Results

In this section, we report the results of the experiments described in the previous section. These results support our central argument: Heterogeneous response intensity ranges improve swarm performance, when agents respond probabilistically, due to increased inter-agent variability and the decoupling of response threshold values and response intensity values. In addition, our results support those of previous work in demonstrating that response probabilities $p < 1.0$ allow swarms to recover more quickly from agent extinctions.

The data support the following performance improvements for heterogeneous intensity ranges, relative to homogeneous intensity ranges, for lower response probabilities and paths with frequently changing task demands:

– reduced average positional difference between the target and the tracker
– reduced variability, within a run, in average positional difference between the target and the tracker
– reduced difference between the target and tracker path lengths
– more accurate target path tracking
– reduced task switching

Figure 2 illustrates the effect of heterogeneous intensity ranges on target tracking when response probability is reduced. The top row shows that when response probability $p = 1.0$, homogeneous and heterogeneous intensity ranges

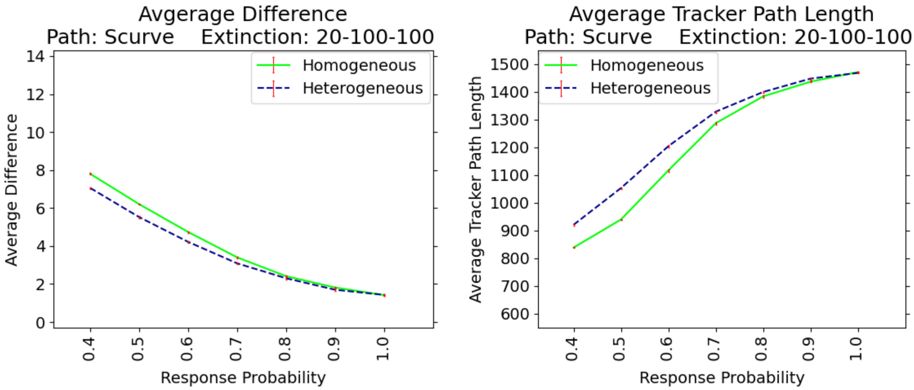

Fig. 4. Average positional difference and tracker path length for homogeneous and heterogeneous intensity ranges for target path s-curve for 100 runs. Error bars are shown in red. Both quantities are improved for heterogeneous ranges. (Color figure online)

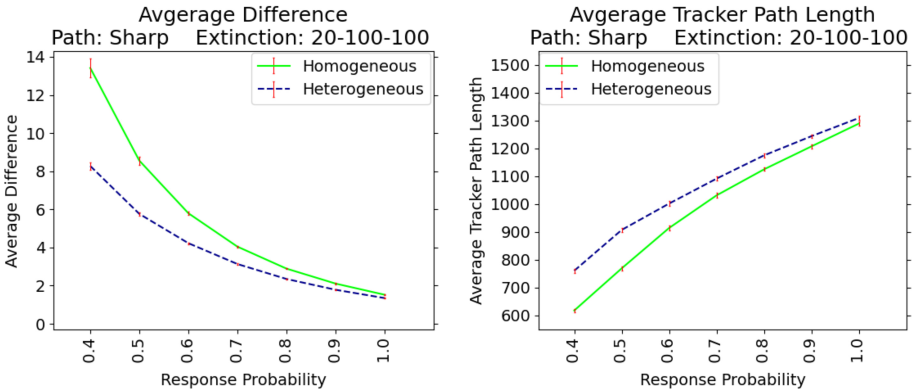

Fig. 5. Average positional difference and tracker path length for homogeneous and heterogeneous intensity ranges for target path sharp. Error bars are shown in red. Both quantities are improved for heterogeneous ranges. (Colour figure online)

produce similar results, with the tracker (red) staying close to the target (blue) throughout the run. Note that there is minimal degradation of performance as agents are killed at 100 time step intervals. In the bottom row, the response probability $p = 0.6$. Thus, there is probability 0.4 that an agent fails to act when its response threshold is met. As a consequence, system performance suffers – recall that parameter values are optimized for $p = 1.0$. We note that tracking is significantly better for heterogeneous intensity ranges than for homogeneous ranges when $p = 0.6$.

Figures 3, 4, and 5 provide data for the tracking effects observed in Fig. 2 for paths random, s-curve, and sharp, respectively. In each figure, the left plot shows

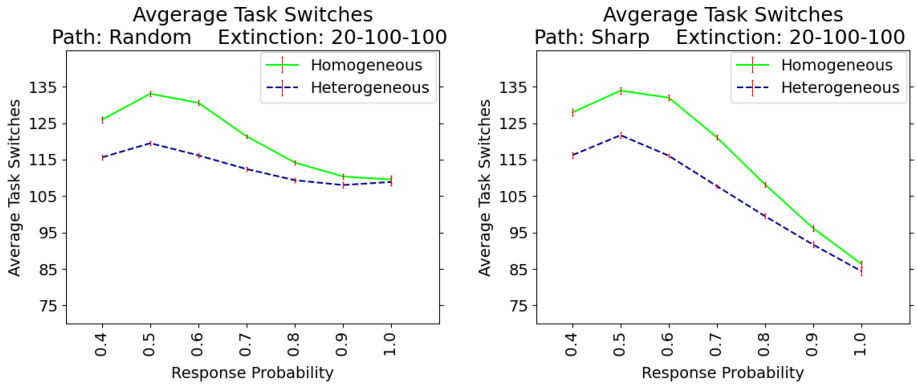

Fig. 6. Average task switches, per agent, for homogeneous and heterogeneous intensity ranges, for target paths random and sharp. These paths have frequently changing task demands. In both cases, heterogeneous ranges reduce the number of switches.

average positional difference between target and tracker for response probabilities $p \in [0.4, 1.0]$ for both homogeneous and heterogeneous response intensity ranges. The right plot shows average tracker path length for the same response probabilities and response intensity ranges. Recall that target path length is 1500. Each data point represents 100 runs. 95% confidence intervals, though quite small in most cases, are shown in red. The data show that at lower response probabilities, the average difference is lower and tracker path length is closer to target path length for runs with heterogeneous response intensity ranges than for runs with homogeneous response intensity ranges.

Figure 6 illustrates the effect of heterogeneous intensity ranges on the average number of switches per agent for paths random (left) and sharp (right). Heterogeneous intensity ranges allow the swarm to perform fewer task switches, particularly when agents respond probabilistically. At response probability $p = 0.6$, the difference is approximately 15 fewer task switches per agent or 3000 fewer switches during a run for a population of 200 agents. Because task switches can incur a cost in real-world applications, this is a significant improvement.

The observed trends are explained as follows. With homogeneous intensity ranges, all frequent actors for a task have similar response intensity values due to the common maximum value. As frequent activation results from low response thresholds, this results in a coupling of the two values: small $\theta \rightarrow$ large γ. In contrast, heterogeneous intensity ranges have different sizes and different minimum and maximum values. The smallest range size $d = 0.6$ and the smallest possible $\gamma_{i,D}min = 0.3$ resulting in a minimum intensity range of $[0.3, 0.9]$. The largest possible $\gamma_{i,D}max = 2.4$. As with homogeneous intensity ranges, frequent actors may reach the maximum intensity value in their range, however, these maximum values vary considerably making the swarm in general, and the group of frequent actors in particular, more diverse. In this way, $\gamma_{i,D}$ is far less dependent on $\theta_{i,D}$. Thus, heterogeneous intensity ranges decouple response intensity values from response threshold values.

This decoupling has multiple effects. First, it allows the swarm to better adapt to frequently changing task demands. This occurs because when task demands change frequently, agents are unlikely to maximize their intensity values through activation. This may result in insufficient intensities, for those agents that activate due to low thresholds, to maintain a small positional difference with the target. The greater diversity of intensity ranges can mitigate this. Second, it helps regulate swarm behavior, in the short-term, after an agent extinction because survivors – agents with higher response thresholds – may have higher response intensities than is possible with homogeneous intensity ranges. Thus, the swarm is better able to meet task demand. Of course, the random nature of intensity range creation could result in a swarm with too few agents with high intensity ranges but due to the population size used, this is unlikely.

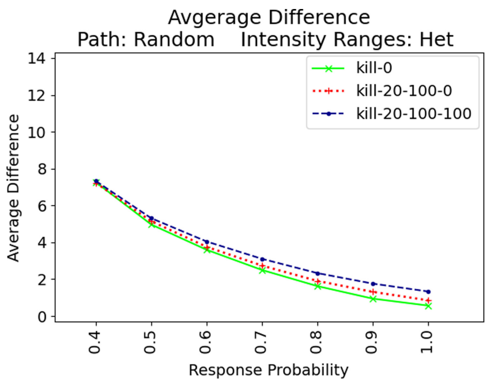

Fig. 7. Average positional difference heterogeneous intensity ranges with each agent extinction implementation for target path random. This demonstrates that extinction type does not significantly affect the swarm's ability to track the target.

The results presented above are for runs using extinction `kill-20-100-100` in which 20 agents are killed every 100 time steps. Extinction types `kill-0` and `kill-20-100-0` are also used in our experiments. Figure 7 illustrates why we focus the discussion on a single extinction type. The figure shows that average positional difference does not vary significantly with changes in extinction. The same trend is observed for average path length. Therefore, we choose to concentrate the analysis on `kill-20-100-100` to simplify the presentation. The y-axis range in Fig. 7 is the same as in Figs. 3, 4, and 5 to facilitate comparison.

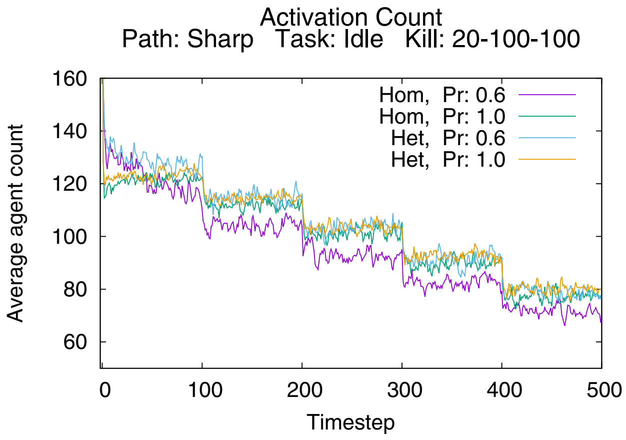

Fig. 8. Idle agent counts for target path sharp. Homeogeneous and heterogeneous response intensity ranges are compared for response probabilities 0.6 and 1.0. For response probability 0.6, heterogeneous intensity ranges result in an increase in idle agents, reducing costs for the swarm. (Color figure online)

Figure 8 shows an additional effect of heterogeneous intensity ranges. With reduced response probability, the number of idle agents (blue line) is greater than when intensity ranges are homogeneous (purple line). Recall that agent activation has costs that can include fuel and wear on the agents. Thus, a higher number of idle agents is desirable. Note that the number of idle agents decreased through the runs represented in the figure due to the reduction in the number of agents through extinctions.

5 Conclusions and Future Work

In this work we explore the effects of a little-studied and promising form of inter-agent variation: response intensity. Expanding on previous work that shows the benefit of heterogeneous response intensity values that vary within a homogeneous range, we implement response intensity values that vary within heterogeneous response intensity ranges. Our system also uses homogeneous response probability and heterogeneous response thresholds.

We find that heterogeneous response intensity ranges provide significant improvement over homogeneous response intensity ranges for decreased response probabilities and problems with frequently changing task demands for a 2-D tracking problem. The improvement is seen in all measures of swarm performance: average positional difference, average tracker path length, and average number of task switches. The observed improvements are due to the decoupling effect that heterogeneous intensity ranges have on response intensity values and response probability values. This results in far more diversity among frequent actors and the backup agents that replace them when agent extinctions occur.

In future work, we will test our model on a more complex task allocation problem and explore additional forms of inter-agent variation. In addition, we plan to investigate heuristic methods for initializing the values of response thresholds and response intensities.

Acknowledgement. This work is supported by the National Science Foundation under grant IIS1816777.

References

1. Bonabeau, E., Theraulaz, G., Deneubourg, J.L.: Quantitative study of the fixed threshold model for the regulation of division of labor in insect societies. In: Proceedings: Biological Sciences, pp. 1565–1569 (1996)
2. Bonabeau, E., Theraulaz, G., Deneubourg, J.L.: Fixed response thresholds and the regulation of division of labor in insect societies. Bull. Math. Biol. **60**, 753–807 (1998). https://doi.org/10.1006/bulm.1998.0041
3. Dornhaus, A., Holley, J., G.Pook, V., Worswick, G., Franks, N.R.: Why do not all workers work? Colony size and workload during emigrations in the ant temnothorax albipennis. Behav. Ecol. Sociobiol. **63**, 43–51 (2008). https://doi.org/10.1007/s00265-008-0634-0
4. Hackwood, S., Beni, G.: Self-organization of sensors for swarm intelligence. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 819–829 (1992)
5. Jeanne, R.L.: Regulation of nest construction behavior in Polybia occidentalis. Animal Behav. **52**, 473–488 (1996)
6. Jones, J.C., Myerscough, M.R., Graham, S., Oldroyd, B.P.: Honey bee nest thermoregulation: diversity promotes stability. Science **305**(5682), 402–404 (2004)
7. Krieger, M.J.B., Billeter, J.B.: The call of duty: self-organised task allocation in a population of up to twelve mobile robots. Robot. Auton. Syst. **30**, 65–84 (2000)
8. Krieger, M.J.B., Billeter, J.B., Keller, L.: Ant-like task allocation and recruitment in cooperative robots. Nature **406**, 992–995 (2000)
9. Langridge, E.A., Franks, N.R., Sendova-Franks, A.B.: Improvement in collective performance with experience in ants. Behav. Ecol. Sociobiol. **56**, 523–529 (2004). https://doi.org/10.1007/s00265-004-0824-3
10. Mathias, H.D., Wu, A.S., Ruetten, L., Coursin, E.: Improving multi-agent system coordination via intensity variation. In: Proceedings of the 33rd International Florida Artificial Intelligence Research Society Conference (2020)
11. Oster, G.F., Wilson, E.O.: Caste and Ecology in the Social Insects. Princeton University Press, Princeton (1978)
12. Ravary, F., Lecoutey, E., Kaminski, G., Chaline, N., Jaisson, P.: Individual experience alone can generate lasting division of labor in ants. Curr. Biol. **17**, 1308–1312 (2007)
13. Riggs, C., Wu, A.S.: Variation as an element in multi-agent control for target tracking. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 834–841 (2012)
14. Theraulaz, G., Goss, S., Gervet, J., Deneubourg, J.L.: Task differentiation in Polistes wasp colonies: a model for self-organizing groups of robots. In: Proceedings of the 1st International Conference on Simulation of Adaptive Behavior: From Animals to Animats, pp. 346–355 (1991)

15. Weidenmüller, A.: The control of nest climate in bumblebee (*Bombus terrestris*) colonies: Interindividual variability and self reinforcement in fanning response. Behav. Ecol. **15**, 120–128 (2004)
16. Wu, A.S., Mathias, H.D., Giordano, J., Hevia, A.: Effects of response threshold distribution on dynamic division of labor in decentralized swarms. In: Proceedings of the 33rd International Florida Artificial Intelligence Research Society Conference (2020)
17. Wu, A.S., Wiegand, R.P., Pradhan, R.: Using response probability to build system redundancy in multi-agent systems. In: Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems, pp. 1343–1344 (2013)
18. Wu, A.S., Wiegand, R.P., Pradhan, R.: Building redundancy in multi-agent systems using probabilistic action. In: Proceedings of the 29th International Florida Artificial Intelligence Research Society Conference. pp. 404–409 (2016)
19. Wu, A.S., Wiegand, R.P., Pradhan, R.: Response probability enhances robustness in decentralized threshold-based robotic swarms. Swarm Intell. (2020). https://doi.org/10.1007/s11721-020-00182-2
20. Wu, A.S., Wiegand, R.P., Pradhan, R., Anil, G.: The effects of inter-agent variation on developing stable and robust teams. In: Proceedings of the AAAI 2012 Spring Symposium: AI, The Fundamental Social Aggregation Challenge, and the Autonomy of Hybrid Agent Groups (2012)