Effects of Response Threshold Distribution on Dynamic Division of Labor in Decentralized Swarms

Annie S. Wu[1], H. David Mathias[2], Joseph P. Giordano[1], and Anthony Hevia[1]
[1]University of Central Florida, Orlando, Florida 32816, aswu@cs.ucf.edu
[2]University of Wisconsin - La Crosse, La Crosse, WI 54601, dmathias@uwlax.edu

Abstract

In this paper, we investigate how the distribution of response threshold values affects the ability of decentralized swarms to dynamically achieve appropriate division of labor in response to changing task demands. Inter-agent variation of response thresholds is a common method for de-synchronizing decentralized agents, which can result in more effective division of labor. We present a systematic study of three different distributions that are relevant to natural and artificial swarms. We use each of these distributions to generate the agent response thresholds in a swarm and examine the accuracy and stability of the swarm's performance on a collective control problem.

Introduction

We investigate how the distribution of response threshold values affects the ability of decentralized swarms to dynamically achieve appropriate division of labor in response to changing task demands. We focus on problems in which task stimuli are globally available to all agents in the swarm. The response threshold method for task allocation is a commonly used approach for generating division of labor in decentralized swarms. This method is a reactive method in which agents dynamically decide which task to respond to based on the task stimuli sensed at any give time. Decentralized swarms are robust because the lack of a central controller means that there is not a single point of failure. The lack of a central controller, however, also makes it more difficult to coordinate agents such that the group as a whole responds intelligently and efficiently to multiple task demands.

For a multi-agent system (MAS) to be able to respond intelligently to different situations or states, the individual agents that make up the system must be able to respond differently to the same input state (Ashby 1958). For problems in which task stimuli are sensed locally, such as deciding what resource to retrieve based on locally observed distributions (Jones and Mataric 2003; Lee and Kim 2017; Lerman et al. 2006), it is easier for decentralized agents to distribute themselves among different tasks because agents are likely to sense different stimuli at any given time and, thus, respond differently. For problems in which task stimuli are somewhat

or completely global, such as deciding whether or not to forage based on the level of a common food store (Castello et al. 2013; Krieger and Billeter 2000), ensuring that all agents do not behave identically can be a challenge as all agents sense the same task stimuli. In such problems, variable agent response is generally accomplished in one of two ways: agents respond probabilistically to task stimuli (Bonabeau, Theraulaz, and Deneubourg 1996; Kalra and Martinoli 2006; Price and Tino 2004) or agents are assigned different thresholds for the same task (Campbell, Riggs, and Wu 2011; Krieger and Billeter 2000; Riggs and Wu 2012). Either of these approaches may be extended such that agent thresholds dynamically adapt over time (Castello et al. 2013; Theraulaz, Bonabeau, and Deneubourg 1998).

We are interested in the second approach, where agents are assigned different thresholds for the same task, for three reasons. First, variable thresholds make some agents more responsive to certain stimuli than other agents. If there is a cost for agents to switch tasks, variable thresholds can reduce such costs because the same (most responsive) agents for a given task are the first to respond and most likely to remain on a task. While probabilistic response can produce variable agent behavior, qualitatively, all agents are still responding to any stimuli in the same way: with the same probabilistic response. Thus, probabilistic response alone provides no mechanism for specialization or reducing task switching. Second, understanding what threshold distributions are best for what situations in static threshold swarms will allow us to more effectively evaluate how well a swarm with dynamically adapted thresholds is performing. Third, while dynamically adapted thresholds theoretically allow a swarm to adjust its threshold distributions to whatever distribution is appropriate for the task demands at any given moment, previous work finds that once a dynamic system has adapted to a set of task demands, it can be difficult for the system to re-adapt to new demands (Kazakova and Wu 2018; Theraulaz, Bonabeau, and Deneubourg 1998). Dynamically adapted thresholds typically use a positive feedback loop which commonly take a system to states that are difficult to subsequently escape. As a result, systems that use dynamically adapted thresholds may not as adaptable as expected, and they may be biased by their first task. (Meyer et al. 2015)

provide an example problem where static variable thresholds are more effective than dynamically adapted thresholds.

Studies on both natural (Jones et al. 2004; Weidenmüller 2004) and artificial (Krieger and Billeter 2000; Riggs and Wu 2012) decentralized swarms have established that inter-agent variation of response thresholds can result in more stable and effective division of labor. To our knowledge, however, there is little work studying how different threshold distributions affect swarm behavior (Campbell, Riggs, and Wu 2011). The distribution of threshold values among agents determines the rate at which agents enter and leave the workforce, which can potentially impact the responsiveness and stability of the swarm. As a result, the choice of distribution to use in an artificial swarm may have a significant effect on how the swarm responds to changes in task demand.

We perform a comparative study on how three distributions of response thresholds affect swarm behavior on problems with both gradually and abruptly changing task demands. We perform this study on a collective tracking problem which allows us to systematically define problems with gradually changing and abruptly changing task demands. Specific questions that we investigate are:

- Does the distribution of agent thresholds affect how well the system responds to task needs?
- Are different distributions better for different situations? Situations of interest include gradually changing versus abruptly changing task demand.

Our model

The testbed that we use is a collective tracking problem. The problem consists of a target that moves in various prescribed paths in a 2D space and a tracker that is controlled collectively by a swarm of decentralized non-communicating agents. In each timestep, the target moves a fixed distance in a direction that is specified by the selected path, and the agents in the swarm attempt to collectively move the tracker to match the target's movement. Agents can choose from one of four possible tasks – *PUSH_NORTH*, *PUSH_EAST*, *PUSH_SOUTH*, *PUSH_WEST* – or remain *IDLE*, and the actions of all agents in the swarm are aggregated to determine the tracker movement in that timestep. As a result, the swarm's goal is to ensure that an appropriate number of agents are allocated to each task at any given time. Each run consists of a fixed number of timesteps or moves. Over the course of a run, we measure the average and maximum distance between the target and tracker, the lengths of the paths each travels, and the number of times agents switch tasks.

The authors recognize and concur that there are more efficient and effective ways to achieve tracking. We use this problem as our testbed because it defines a clear set of tasks on which agents must self-organize, and the target paths provide a systematic way to define and study dynamically changing task demands.

Target movement creates dynamic task demands

In each run, the target moves at a fixed speed of S_{target} units per timestep along a selected path. The movement of the target and the resulting relative positions of the target and tracker determines the task demands perceived by the swarm in each timestep. For example, if the target is due east of the tracker, the swarm will detect a non-zero task demand for pushing east and zero demand for all other directions. If the target is northwest of the tracker, the swarm will detect non-zero task demands for pushing north and west and zero task demand for south and west. The distance between the target and tracker determines the magnitude of the task demands.

A separate task demand value is calculated for each of the four directions by subtracting the tracker position from the target position. In the north and south directions, this is the difference between the target's and tracker's y-coordinate values; in the east and west directions, the x-coordinate values. Negative differences are set to zero.

In order to examine different task demand scenarios, we test on target paths with both gradually and abruptly changing task demands. The former are paths in which the target heading changes in small increments from one timestep to the next. These paths include four serpentine paths with increasing period and amplitude (s-curve10, s-curve20, s-curve30, s-curve40) and one path in which the target makes small random changes in its heading in each timestep (random). The latter are paths in which the target heading may change by a large amount. These paths include four sawtooth paths with increasing period and amplitude (zigzag10, zigzag20, zigzag30, zigzag40) and one path in which the target probabilistically changes its heading to a random new heading in each timestep (sharp).

Tracker movement

The maximum distance that the tracker can move in one timestep is $S_{Max} = Ratio \times S_{target}$ where $Ratio \geq 1.0$. Thus, the tracker's maximum speed is as fast or faster than the target speed.

Each agent has a unique threshold for each of the four tasks that it can choose. Thresholds are assigned at the start of each run and remain fixed throughout the run. When the task demand exceeds an agent's threshold for that task, the agent will consider acting on that task. If only one task exceeds an agent's threshold, that agent will choose to act on that task. If more than one task exceeds an agent's corresponding thresholds, the agent randomly selects one of those tasks. If no tasks exceed an agent's threshold, the agent will remain idle. Each agent can act on at most one task at a time.

Agent decisions are aggregated by calculating the percentage, P_D, of total agents that select to push in each direction, D. All non-zero P_D values are summed and the resulting vector multiplied by S_{Max} to generate the tracker's new position. For this problem, there can be at most two non-zero P_D values in any timestep.

Threshold distributions

We are specifically interested in how the distribution of threshold values among the agents of a swarm affect the swarm's ability to satisfy changing task demands. We examine three distributions – uniform, Gaussian, and Poisson – and compare them with a system in which all agents have the same homogeneous constant threshold. Uniform is a com-
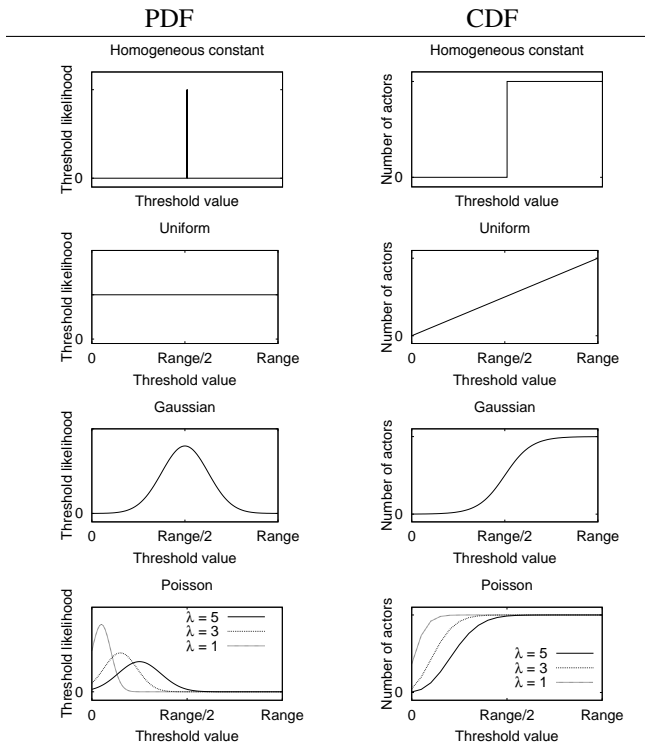
	PDF	CDF

Table 1: Threshold distributions tested.

Parameter	Value
Population size	200
Time steps	500
Target step length, S_{Target}	3
Tracker step multiplier, $Ratio$	3
Range	10

Table 2: Tracking simulation parameter settings
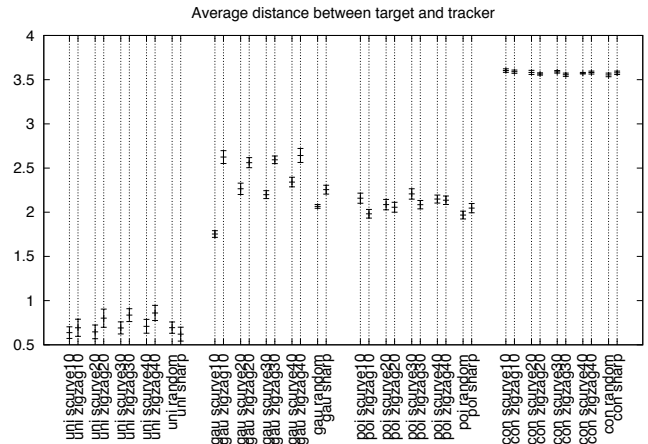

Figure 1: Average distance between target and tracker over a run, averaged over 20 runs with 95% confidence intervals.

monly used distribution in artificial systems. Gaussian and Poisson are relevant in natural systems (Frank 2009).

Table 1 shows the distributions that we test. The probability distribution function (PDF) gives, for each threshold value within the range of possible thresholds, the relative likelihood that an agent will be assigned a threshold of each value. The cumulative distribution function (CDF) gives, for each possible threshold value, the expected number of agents that will act at that threshold value. If all agents have the same threshold for a task (top row of Table 1), all agents will respond identically to the task stimulus and the swarm as a whole has only two possible responses to the task: all agents act or no agents act. Variable thresholds stagger the agent response to a task stimulus, resulting in a more gradual system response. $Range$ is a system parameter that specifies the target-to-tracker distance at which all agents will act.

The three variable distributions that we study have noticeably different CDFs which means that agents are entering and leaving the workforce at very different rates. With uniform distribution, the number of agents that act increases linearly as the threshold value increases. With Gaussian distribution, agents initially enter the workforce slowly; the system becomes more responsive at mid-level thresholds; and the rate of entry decreases again at high thresholds. The CDF of the Poisson distribution looks similar to that of Gaussian, but with more immediate responsiveness, which potentially allows it to be more responsive in situations with abruptly changing task demands.

Experimental results

We apply swarms with each of the four distributions from Table 1 to each of the ten target paths and evaluate the swarm's ability to track each target path. For the Poisson distribution, we use $\lambda = 5$. Each experiment consists of 20 runs. Unless otherwise specified, all measurements are averaged over 20 runs and the average and 95% confidence intervals are reported. Table 2 lists the tracking simulation parameter settings used for the experiments presented here. For each experiment, we measure system accuracy by examining the average distance between the target and tracker throughout the entire run and the length of the target and tracker paths. We measure stability by examining the average number of times agents switch tasks during a run.

Accuracy

Figure 1 plots, for each experiment, the average distance between the target and tracker over the duration of a run. Distance is a direct measure of how well a swarm's task allocation (tracker movement) meets the task demand (target movement) in each timestep. The x-axis plots the experiment and the y-axis plots distance. Each experiment is defined by a threshold distribution and a target path. Figure 2 shows an example run from each threshold distribution for s-curve10, zigzag10, s-curve30, and zigzag30.

The results indicate that swarms initialized with a uniform threshold distribution and swarms initialized with a Gaussian threshold distribution are better at tracking gradually changing task demands than abruptly changing task
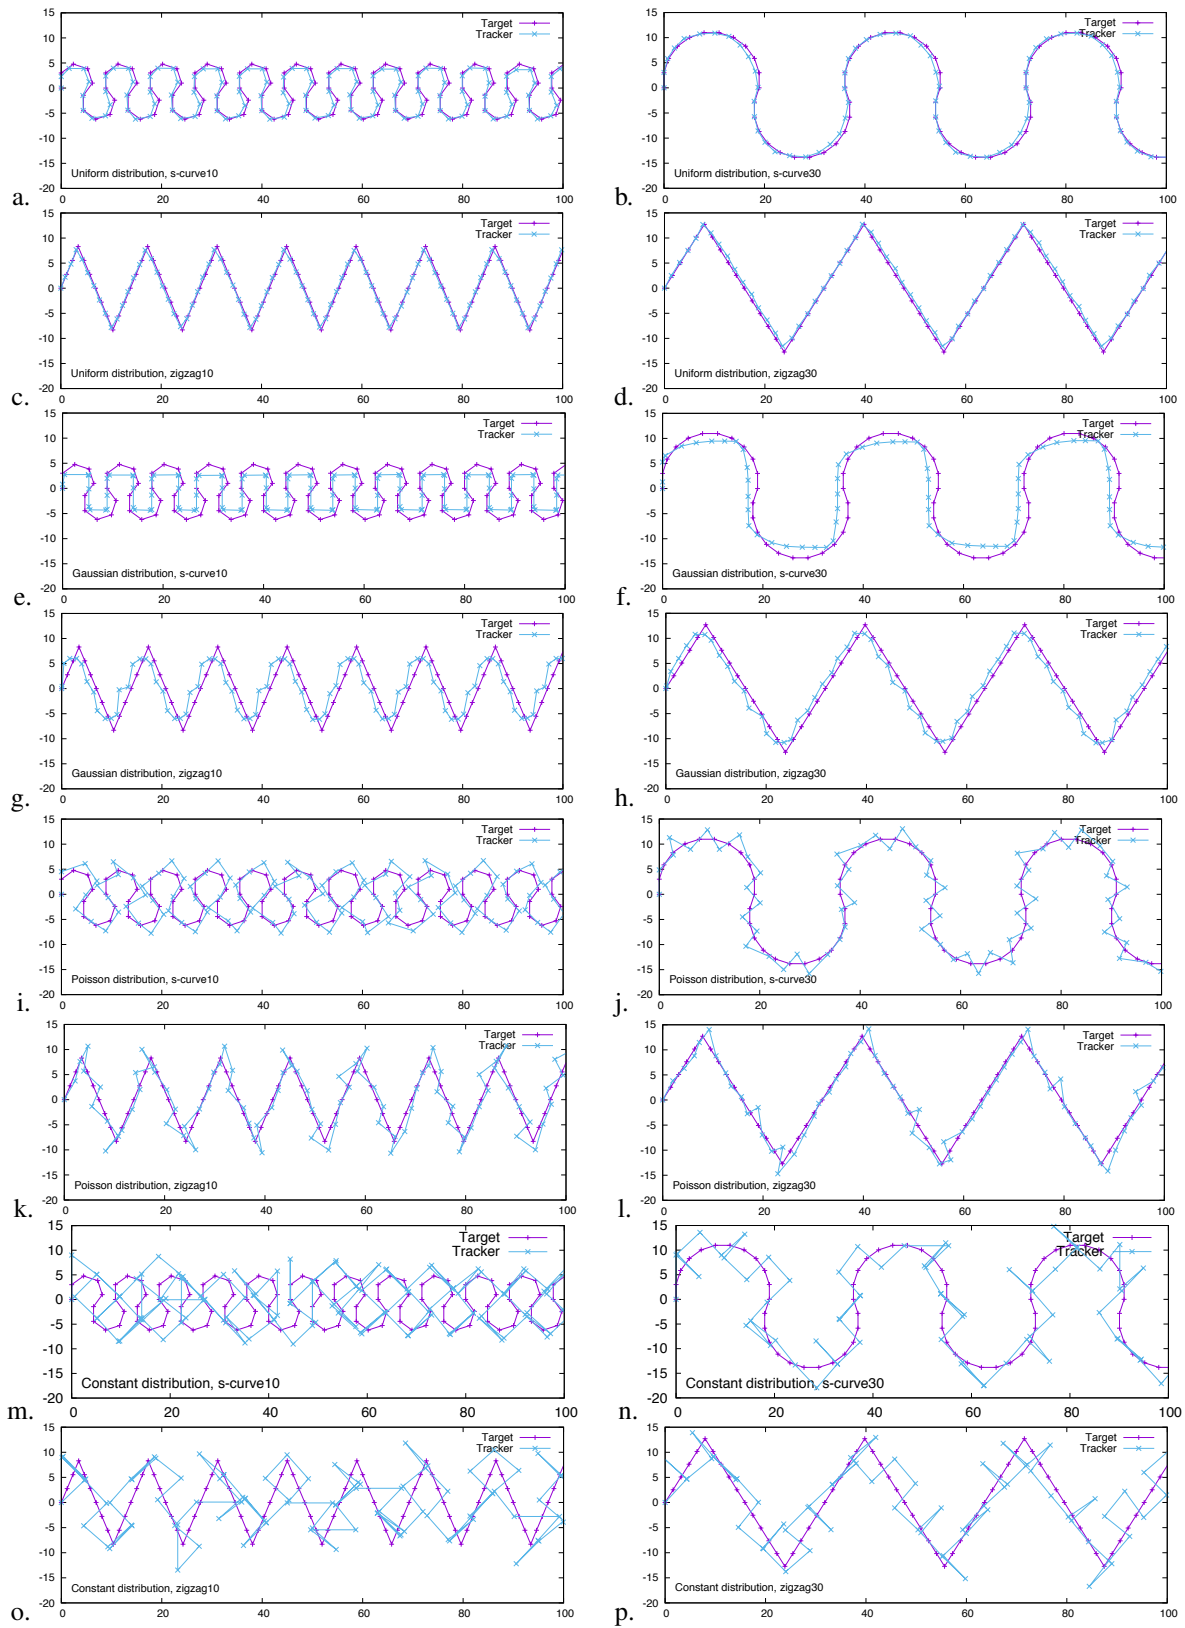
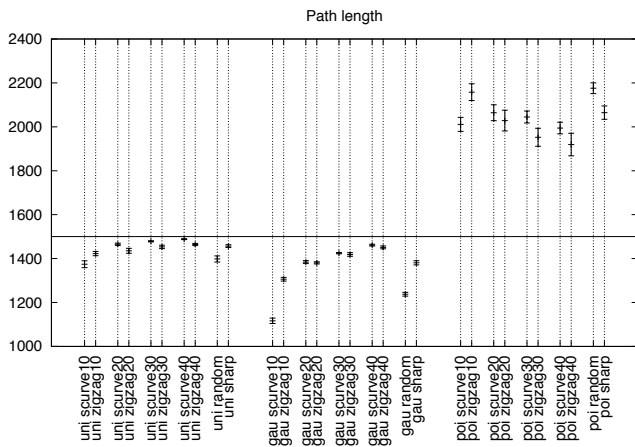Figure 2: Examples of individual runs on s-curve10, zigzag10, s-curve30, and zigzag30.

Figure 3: Tracker path length, averaged over 20 runs with 95% confidence intervals. The optimal path length is 1500.

Figure 4: Average number of task switches over an entire run, averaged over 20 runs with 95% confidence intervals.

demands. Both sets of swarms maintain a shorter distance between the target and tracker on the s-curve (gradual) path than zigzag (abrupt). The difference is significant in Gaussian swarms; not clearly significant in uniform swarms. The Poisson swarms and constant swarms perform slightly better on abrupt over gradual changes. This difference may be due to the fact that the Poisson and constant distributions cause swarms to mount a larger response to smaller task demands than Gaussian or uniform. Examination of Figures 2i-2l shows that Poisson's over-response makes it hard for the swarm to adjust to the small changes in heading that occur in the s-curve paths, and is less of a liability on straight lines of which zigzag is primarily composed. In comparison, constant swarms do badly on all shown examples.

Overall, uniform swarms maintain significantly shorter distances between the target and tracker than the other three distributions. Gaussian and Poisson swarms perform comparably and constant swarms clearly exhibit the worst performance. Constant distribution's performance is even more of an outlier on the remaining evaluation metrics. As a result, to make the plots more readable, we omit the data from the constant swarm experiments from future plots.

Figure 3 shows the tracker's path length for each experiment, averaged over 20 runs. The optimal path length is 1500, as indicated by the horizontal line. For the uniform and Gaussian distributions, tracker path length better matches target path length on s-curve20/30/40 than zigzag20/30/40. The effect is reversed for the Poisson distribution. For all three distributions, the results for s-curve10/zigzag10 are the reverse of the results on the other three paths, and the results on the sharp path are significantly better than that on the random path. The plots in the left column of Figure 2 suggest that the difficulty with s-curve10, for uniform and Gaussian swarms, stems from its small size. Its curves are too tight for the system to accurately follow the target path. Poisson swarms' tendency to over react to small task demands cause it to perform badly on both s-curve and zigzag. The straight lines of the zigzag paths can ameliorate some of this effect,
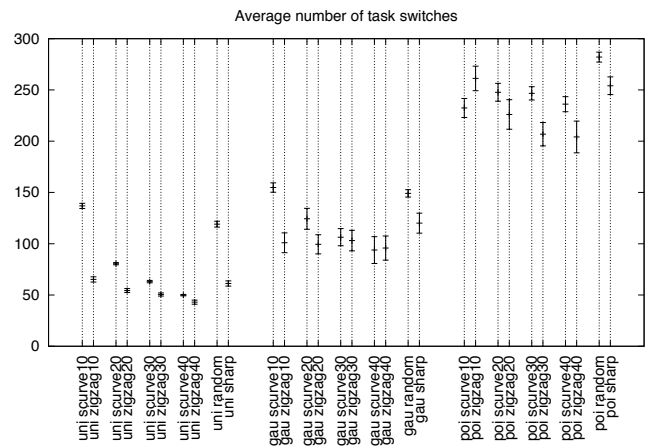
but only if there are significant stretches of straight lines.

Overall, uniform swarms generate tracker path lengths that are most similar to the target path length and Gaussian swarms are a close second. Both uniform and Gaussian swarms generate tracker path lengths that are shorter than the target path length while Poisson swarms generate tracker path lengths that are longer than the target path length. This difference is consistent with what would be expected from the shapes of their PDF and CDF. Uniform and Gaussian distributions increase gradually at low thresholds which is likely to cause them to cut corners, resulting in shorter than optimal paths. The Poisson distribution increases quickly at low threshold values which makes such swarms more likely to overshoot their goals and hence result in a longer than optimal tracker path length. In all three cases, path following accuracy improves as the target path amplitude and period increase. This improvement is likely due to a combination of a smoother path and the ability of the tracker to take smaller steps relative to target path's period and amplitude. Constant swarms are again universally bad, generating paths that are significantly longer than the optimal path.

Stability

Figure 4 plots, for each experiment, the average number of times, over an entire run, that agents switched tasks, averaged over 20 runs. In all comparisons except for the Poisson swarm on s-curve10 and zigzag10, agents execute fewer task switches on zigzag paths than s-curve paths. The preference for zigzag is likely due to the fact that zigzag has more straight lines than s-curve. Once an appropriate division of labor is found, straight lines can be traversed with no task switches indefinitely. In the exceptional case, it appears that the amplitude of zigzag10 is too short for the straight segments to cancel out any overshooting errors made on the corners, resulting in better performance on s-curve10 instead. We also observe a trend of fewer task switches as path size increases. The larger paths have longer straightaways which require fewer task switches to trace.

Overall, uniform swarms are again the best performers with Gaussian a close second. Agents in Poisson swarms execute about twice as many task switches per run as the other two distributions. Agents in constant swarms execute even more task switches, on average, over 400 per run.

Conclusions

In this work, we investigate the impact of four different distributions of response threshold values on the ability of a decentralized swarm to self-organize in response to dynamic task demands. This study is conducted on a collective tracking problem in which agents attempt to track a moving target by aggregating their actions. We examine the performance of uniform, Gaussian, Poisson, and constant swarms on both gradually changing and abruptly changing paths.

Results find that there are explainable differences in how swarms with different threshold distributions perform on different types of paths or problems. In terms of the path traveled by a tracker, most distributions are better able to minimize target to tracker distance, and meet the task demand in each timestep, when the target path has gradually changing task demands. Swarms also have more accurate tracker path length with gradually changing demands than with abruptly changing demands. Only in situations where task demands change rapidly, is the advantage of gradual over abrupt lost. Most distributions experience fewer task switches on the zigzag paths, despite their abrupt changes in task demands, because the periods in between the abrupt changes consist of straight paths which require fewer task switches.

Among the distributions tested, swarms with a uniform distribution perform best in all evaluation metrics. They maintain a significantly shorter target to tracker distance and are comparable to Gaussian with respect to path length and number of task switches. Swarms with a Gaussian distribution are a close second, performing worse on target to tracker distance. Swarms with a Poisson distribution maintain a similar target to tracker distance as Gaussian, but at path lengths that are significantly longer than optimal. This observation, combined with a high rate of task switches, suggests that, although the Poisson distribution does allow a swarm to meet task demands when measured over an extended period, it does so with a significant amount of wasted energy.

Our results suggest that the type of threshold distribution appears to be a greater distinguishing factor in a swarm's ability to dynamically adapt division of labor than the problem on which the swarm is working. Uniform distribution's consistently strong performance on multiple problem types and the simplicity of static distributions make uniform a strong contender over other static distributions, and possibly over more complex dynamically adapted threshold methods. In spite of the fact that natural systems are often associated with more complex distributions, our results suggest that a simple uniform distribution, which is easy to generate, is not only sufficient but the best choice in computational swarms.

Acknowledgements

References

Ashby, W. R. 1958. Requisite variety and its implications for the control of complex systems. *Cybernetica* 1(2):83–99.

Bonabeau, E.; Theraulaz, G.; and Deneubourg, J.-L. 1996. Quantitiate study of the fixed threshold model for the regulation of division of labor in insect societies. *Proc. Royal Society of London: Biological Sciences* 263(1376):1565–1569.

Campbell, A.; Riggs, C.; and Wu, A. S. 2011. On the impact of variation on self-organizing systems. In *Proc. 5th IEEE Int'l Conf. Self-Adaptive and Self-Organizing Systems*.

Castello, E.; Yamamoto, T.; Nakamura, Y.; and Ishiguro, H. 2013. Task allocation for a robotic swarm based on an adaptive response threshold model. In *Proc. 13th Int'l Conf. Control, Automation, and Systems*, 259–266.

Frank, S. A. 2009. The common patterns of nature. *Journal of Evolutionary Biology* 22(8):1563–1585.

Jones, C., and Mataric, M. J. 2003. Adaptive division of labor in large-scale minimalist multi-robot systems. In *Proc. IEEE/RSJ Int'l Conf Intell Robotics & Systems*, 1969–1974.

Jones, J. C.; Myerscough, M. R.; Graham, S.; and Oldroyd, B. P. 2004. Honey bee nest thermoregulation: Diversity promotes stability. *Science* 305(5682):402–404.

Kalra, N., and Martinoli, A. 2006. A comparative study of market-based and threshold-based task allocation. In *Distributed Autonomous Robotics Systems 7*, 91–101.

Kazakova, V. A., and Wu, A. S. 2018. Specialization vs. re-specialization: Effects of Hebbian learning in a dynamic environment. In *Proc 31st FLAIRS*, 354–359.

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* 30:65–84.

Lee, W., and Kim, D. 2017. History-based response threshold model for division of labor in multi-agent systems. *Sensors* 17(6):1232.

Lerman, K.; Jones, C.; Galstyan, A.; and Mataric, M. 2006. Analysis of dynamic task allocation in multi-robot systems. *Int'l Journal of Robotics Research* 25:225–241.

Meyer, B.; Weidenmuller, A.; Chen, R.; and Garcia, J. 2015. Collective homeostasis and time-resolved models of self-organised task allocation. In *Proc. 9th EAI Int'l Conf Bio-inspired Info & Comm Tech*, 469–478.

Price, R., and Tino, P. 2004. Evaluation of adaptive nature inspired task allocation against alternative decentralised multiagent strategies. In *Proceedings of the Parallel Problem Solving from Nature, LNCS 3242*, 982–990.

Riggs, C., and Wu, A. S. 2012. Variation as an element in multi-agent control for target tracking. In *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems*, 834–841.

Theraulaz, G.; Bonabeau, E.; and Deneubourg, J. 1998. Response threshold reinforcement and division of labour in insect societies. *Proc. Royal Society B* 265:327–332.

Weidenmüller, A. 2004. The control of nest climate in bumblebee (*Bombus terrestris*) colonies: Interindividual variability and self reinforcement in fanning response. *Behavioral Ecology* 15:120–128.