

Exploiting Intrinsic Multi-Agent Heterogeneity for Spatial Interference Reduction in an Idealised Foraging Task

Chris Bennett¹, Jonathan Lawry¹ and Seth Bullock²

¹ Department of Engineering Maths, University of Bristol, Bristol, UK

² Department of Computer Science, University of Bristol, Bristol, UK
christopher.bennett@bristol.ac.uk

Abstract

Typically, collective behaviour research has tended to focus on behaviour arising in populations of *homogeneous* agents. However, humans, animals, robots and software agents typically exhibit various forms of *heterogeneity*. In natural systems, this heterogeneity has often been associated with improved performance. In this work, we ask whether spatial interference within a population of co-operating mobile agents can be managed effectively via conflict resolution mechanisms that exploit the population's intrinsic heterogeneity. An idealised model of foraging is presented in which a population of simulated ant-like agents is tasked with making as many journeys as possible back and forth along a route that includes tunnels that are wide enough for only one agent. Four conflict resolution schemes are used for determining which agent has priority when two or more meet within a tunnel. These schemes are tested in the context of heterogeneous populations of varying size. The findings demonstrate that a conflict resolution mechanism that exploits agent heterogeneity can achieve a significant reduction in the impact of spatial interference. However, whether or not a particular scheme is successful depends on how the heterogeneity that it exploits is implicated in the population-wide dynamics that underpin system-level performance.

Introduction

Alife has explored collective behaviour in populations of simple agents. Often these agents are homogeneous in order to keep models tractable. However, in nature heterogeneity is a key property of many systems that exhibit interesting or sophisticated collective behaviour. For example, biodiversity has been linked to increased stability and productivity of an ecosystem by improving its resilience to environmental effects (Tilman et al., 2006; Hooper et al., 2005). Individual differences in a group can improve collective problem solving in humans (Page, 2008) and behavioural differences, together with social connectivity, can change the attack speed of social spiders (Hunt et al., 2019).

This paper explores one particular aspect of this wide ranging topic by taking the concepts of functional diversity and fitness and asking if they can be exploited to create a better engineered system. Specifically we consider whether linking the “fitness” of an agent to its priority (status) in the

system produces better overall system performance. How do different schemes for assigning priority affect the results? And, how does increased functional diversity (heterogeneity) in the system change the behaviour?

Asking these questions in this specific case contributes to the wider aim of finding design rules for exploiting the right kind of heterogeneity in multi-agent systems, thereby moving away from the view that regards all heterogeneity as a potential problem in engineered systems.

We begin with a brief discussion of interference in multi-agent systems, followed by an ecological definition of fitness as arising from a combination of extrinsic and intrinsic traits, and consider the hypothesis that giving priority to the fittest agent during local interference interactions will improve system level performance. To test these concepts, four different mechanisms for assigning priority are described. These are then evaluated within a simple simulation environment in which a population of heterogeneous ant-like agents are tasked with repeatedly navigating to a source of food and carrying the food home along a route with narrow tunnels. The findings show that the performance achieved using each mechanism is sensitive to both the environment and the heterogeneous traits of the population.

Interference

Interactions between agents co-operating on a task can be broadly classified as being either beneficial or detrimental to the overall function of the group. Here, we focus on “same place, same time” interference (Goldberg and Matarić, 1997) where two or more agents are not able to access the same limited resource at the same time. This form of interference is present in a wide range of scenarios including organisms foraging for food (Vahl et al., 2005) and robots navigating cluttered environments (Trautman and Krause, 2010). The specific question addressed in this paper is how best to resolve same place, same time interference between two or more agents that may differ in some of their intrinsic characteristics.

One approach is to limit or avoid spatial interference completely by suitably partitioning the environment (and task),

for instance via the use of a hand over area to avoid foraging agents crowding the “nest” (Pini et al., 2009) or dividing the space and assigning a single agent to operate in each spatial zone (Schneider-Fontán and Mataric, 1998). This is comparable to the approach taken by territorial animals (Bullock, 2016). However, avoiding other agents is not always possible, and in these scenarios, a strategy is needed for deciding which agent has priority.

Agents that are obstructing one another could engage in agonistic interactions in order to determine which should back down. However, it is likely to be less costly to resolve the impasse by employing a conflict resolution mechanism that avoids direct combat. One approach to deciding priority is to base it on a trait of the individual which can be attributed to the likely future success of the population. In nature, this trait may be based on physical characteristics such as body size and success described quantitatively as the fitness of the individual (Blanckenhorn, 2000).

When we create artificial systems, for example a team of robots, the concepts of fitness and traits are no less relevant but can be harder to define. When the goal is to maximise team performance then an analogy to functional traits may be appropriate (Petchey and Gaston, 2006). These are a subset of the agent’s traits which affect the performance of the ecosystem. While fitness may not be directly applicable without a mechanism for reproduction, researchers have developed proxies for it in robotic systems by relating a robot’s current state to its contribution to a task. In Mayya et al. (2019), a swarm distribution task algorithm is proposed which is designed to reduce spatial interference by enabling robots in densely packed regions to judge when it is more beneficial to move away and make space than it is to stay and participate in the task. Alternatively, in Brown et al. (2005) the amount an agent has invested in the task is used to assign higher priority to an agent if it has more to lose by giving way to another agent. In both these approaches, the proxy for the fitness of each agent is a function of its current state. All agents are intrinsically identical - they only differ in terms of some transient properties related to their current circumstances or history of interactions with the world.

Contrasting these perspectives from ecology and engineering highlights that a fitness proxy may either be related to intrinsic properties of the individual (weight, height, speed...) or extrinsic properties (task investment, personal space, level of aggression..). Previous work in the multi-agent literature on agent interference has often used physically homogeneous agents. In Vaughan et al. (2000), two homogeneous robots would attempt to pass through the same narrow opening. It was found that a proxy for “aggression” was effective at resolving interference and increased the collective performance of the population. Surprisingly, there was no significant difference between assigning a robot’s “aggression” at random vs a fixed hierarchy. The authors speculated that if a functional difference between the robots




Figure 1: The experimental environment showing two tunnels and five ant-like agents foraging for food. Green agents are currently in the Navigating state, the yellow agent currently is Afraid and the Red agent currently is Brave.

had existed (i.e., the population had been heterogeneous), then some mechanisms for assigning priority, based on a proxy for fitness, would be more beneficial than others.

Here, we build on these prior works via two experiments using heterogeneous populations which test the utility of different intrinsic properties for resolving interference. The hypothesis is that for the purposes of resolving “same place, same time” interference, the overall performance of the group is improved when the conflict resolution is informed by a proxy for fitness that determines an agent’s priority. In the first experiment, speed is taken as the proxy since it directly affects how much food an agent can transport in a given time. In experiment two, the proxy is based on the maximum sensor range within which an ant can detect other ants that it is in conflict with. Through these two experiments, we explore the way in which the nature of some heterogeneity in an agent’s intrinsic properties impacts the efficacy of conflict resolution mechanisms intended to reduce the impact of inter-agent interference.

Experimental Approach

Foraging is a popular scenario in the multi-agent and swarm literature because it can be implemented using simple behaviours for each agent while still creating interesting results relevant to the real-world (Brambilla et al., 2013; Liu and Winfield, 2010; Dugatkin, 2002; Pitonakova et al., 2018). Here, we simulate a small population of ant-like organisms tasked with transporting food from one location to another along a route that includes a narrow tunnel (fig. 1). The virtual ants can pass each other without penalty outside the tunnel, but inside the tunnel interference occurs, preventing one agent from overtaking a slower agent or moving past an agent that is approaching from the opposite direction. The tunnel is only a single lane wide and therefore the ants must decide who will back up and give way to avoid a deadlock. Note that when a faster ant approaches a slower ant and both are travelling in the same direction, both ants continue along their current path at the speed of the slower ant.

If two ants are identical, then it can be presumed they would take the same action when they meet travelling in op-

posite directions. That is, either both back up or both go forward. To avoid a deadlock, there needs to be some symmetry breaking (or diversity) between the agents. In the following we consider two situations in which agents differ in terms of their priority and their functional traits.

Assigning Priority

Four simple mechanisms for assigning priority are as follows:

- **Arbitrary (Transient):** Each agent is assigned a unique random value when it enters a conflict. This value is lost when the conflict is resolved. During a conflict, the agent with the highest value assumes priority.
- **Arbitrary (Fixed):** Each agent is assigned a unique random value at the start of the simulation. In a conflict, the agent with the highest value assumes priority.
- **Fittest First:** Each agent is assigned a value based on a proxy for fitness. In a conflict, the agent with the highest value assumes priority.
- **Fittest Last:** Each agent is assigned a value based on a proxy for fitness. In a conflict, the agent with the *lowest* value assumes priority.

Assigning Functional Traits

From an ecological perspective, a functional trait is a measurable characteristic of an organism that can be linked to the overall function of the ecosystem (Petchey and Gaston, 2006). Here we take a simplified approach and assign each ant a unique set of traits. In experiment one, the trait is speed and there is an implicit assumption that a faster ant will be “fitter” since it has the potential to transport more food over a given time period. In each simulation, speeds were assigned to agents such that the mean population speed remains constant across all scenarios. This avoids skewing the results to favour populations containing agents with a higher average speed. In the idealised case where no spatial interference occurs, this causes the total food collected to increase (linearly) with the number of ants, but the average food collected per ant to remain constant (see fig. 2).

Ecosystem Concepts of Diversity

The experiment uses heterogeneous populations and therefore a workable definition of diversity is needed. The diversity of a population can be defined based on species richness (the number of types present), evenness (how many of each type are present) and divergence (the difference between types) (Mason et al., 2005). An agent type is defined via its traits. In each simulation, each agent has a unique set of traits and therefore richness is varied by changing the number of ants present with evenness remaining constant.

Divergence is a measure of how different two agents are. Metrics that address this for two individuals are often based




Figure 2: Boxplots of the total food transported in experiment 1 with no interference. This is a maximum under idealised conditions. Mean Speed is 50. Speed range is 10. *Left:* total food. *Right:* total food normalised by the number of agents. Note, speeds are chosen such that the average remains constant for all population sizes.

on some form of difference norm. Compound metrics extend this measure of divergence by combining it with richness and evenness, to give a single number representing how diverse (heterogeneous) a population is. A discussion of the different classifications, norms and metrics together with their relative merits is given by Daly et al. (2018) with particular examples given in Balch (2000); Twu et al. (2014). Here, divergence is varied primarily by changing the separation of speed values within a population. For example in a population of 3 agents with velocities $V = \{35, 40, 45\}$, the mean velocity is $\bar{v} = 40$ and the separation is $\Delta v = 5$. If the velocities were $V = \{30, 40, 50\}$, the mean velocity remains $\bar{v} = 40$ and the separation is now $\Delta v = 10$. In the experiments reported here, we keep the number of different speeds equal to the number of agents and manipulate only the range of speeds, keeping the mean speed constant. (fig. 2).

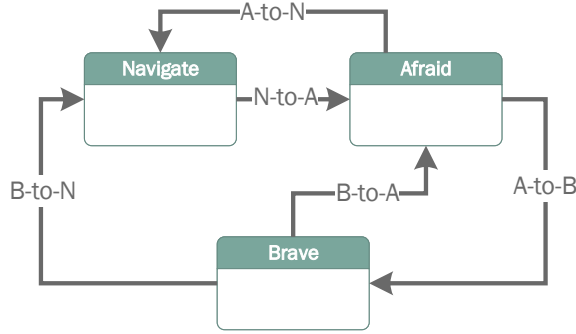
Simulation Environment

The simulation environment was written in python and is available for download from the University of Bristol’s data repository (Bennett et al., 2022)¹.

Finite State Machine A FSM is used to control the behaviour of an ant transitioning between transporting food (Navigating), backing up (Afraid) and advancing with priority over any afraid ants ahead (Brave). The transitions between the states are described by the state machine and table shown in figure 3.

When entering the Afraid state, the ant calculates its priority according to the conflict resolution mechanism being implemented. An ant A maintains a list of other ants that it is currently in conflict with and will transition to Brave if it has the highest priority. New ants are added to A ’s list if they move within conflict range of A , are in the tunnel, and are on a collision course with A . An ant is removed from A ’s conflict list if either 1) if is outside the passive range of A , 2) it is outside the tunnel and A is Brave, or 3) A is Afraid and

¹<https://doi.org/10.5523/bris.3ota45j8tbx332daufr3vp35ad>



Transition	Conditions
N-to-A	not isEmpty(conflict_list)
A-to-N	not in tunnel and isEmpty(conflict_list)
A-to-B	my_priority > max(conflict_list)
B-to-A*	my_priority < max(conflict_list)
B-to-N	isEmpty(conflict_list) or not in tunnel

Figure 3: State machine used by an ant to transport food, back away from another agent when afraid and move forward after winning a contest and becoming brave. *An agent does not re-calculate its priority when it transitions from state Brave to Afraid

more than 20px outside the tunnel. The list is cleared when the agent transitions to the Navigate state.

The condition for transitioning from Afraid to Navigate is designed to ensure an Afraid ant reverses far enough to enable other ants to clear the tunnel and allow a Brave ant to exit. Similarly, the passive range is always larger than the conflict range to prevent an agent being rapidly added to and removed from the conflict list.

Collision Avoidance A simple mechanism is used to prevent two or more ants occupying the same space while in a tunnel. An ant will move at $u(t) \leftarrow max_speed * dt$ provided reaching the position at $p(t + dt) = p(t) + u(t)$ does not collide with, or pass through, another ant. If it does then $u(t) \leftarrow 0$ and the ant doesn't move.

Simulation Parameters

The fixed parameters are given in table 1. The independent variables are the number of ants ([2, 9]), the functional diversity (set to either a speed separation of $\Delta v = 5$ or $\Delta v = 10$) and the scenario for assigning priority.

Conflict Range The visible range of an ant is divided into three segments determining: when ants are visible (Detection Range), when they are close enough to cause a conflict (Conflict Range), and when they are far enough away to be removed from a conflict (Passive Range). If the conflict

Parameter	Value	Units
Time Limit	75000	sec
Step Size (dt)	0.1	sec/tick
Number of repeats	50	–
Detection Range [min,max]	[50,350]	px
Conflict Range [min,max]	[20,350]	px
Passive Range [min,max]	[40,350]	px
Range Growth Rate	+5	px/tick
Range Decay Rate	-2	px/tick
Average Population Speed	50	px/tick
Speed Difference (Δv)	{5,10}	px/tick
Number of Repeats	50	–
Arena Size	1280	pxwide
Agent Size	10	wide px

Table 1: Simulation parameters. Agents move along a fixed line parallel with the x-axis. Measurements in pixels (px). The range of speeds in a population of N is $\Delta v N$.

range is too small then unrecoverable deadlocks can occur (see figure 4). Such deadlocks are resolved by gradually increasing an ant's Conflict, Passive and Detection ranges by +5 for each tick during which the ant is stationary. The ranges then decrease by -2 for each tick during which the ant is moving. A minimum and maximum for these ranges was set and is given in table 1. Agents are initialised at their minimum range values. In the first experiment these limits are the same for the all agents, but in the second experiment each agent has different maximum ranges.

Update Method At $t = 0$, the ants are distributed along the route and assigned a random direction of travel. Updates are performed asynchronously (all agents take a turn to sense and move) and the turn order is randomised at the start of each time step to avoid the possibility of introducing a systematic "first mover advantage".

Step Size The simulation step size was found to have a significant effect on the results with large step sizes (greater than 0.5) creating the possibility of artefacts caused by agents moving too far in a single step. These occurred more frequently when faster agents (and also larger numbers of agents) were present, and they could not be resolved by the combination of priority, state machine and heterogeneity in the population. This is a consequence of the rudimentary collision mechanics which allows a slower agent to move closer to other agents than a faster one. A step size of 0.1 was found to be a good compromise between preventing these artefacts and enabling the simulation to complete in a reasonable time. It is worth stressing that the simulation is designed to investigate heterogeneity and fitness rather than provide an accurate model of robot locomotion.




Figure 4: The conflict range is an important parameter for preventing deadlocks. In this example, the conflict range of 4 and 5 is too small so neither considers the other a source of interference and both agents assume priority.

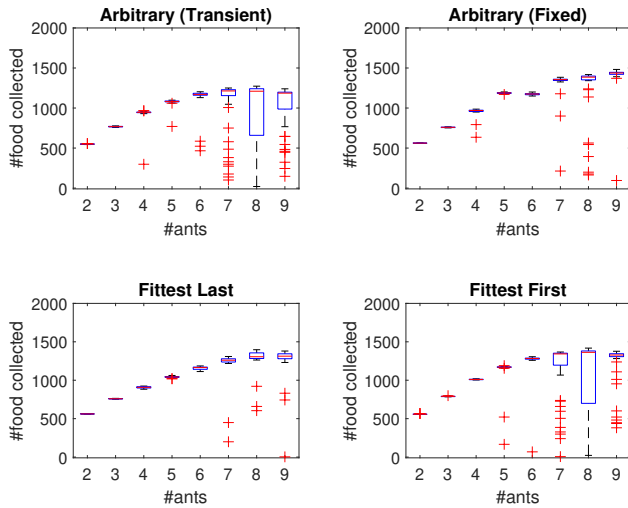


Figure 5: The total food collected for different population sizes with $\Delta v = 5$. The plots show the results from 50 simulation replicates per population size. The heterogeneous trait was speed. Outliers are denoted with a red cross, the box extends from the 25th to 75th percentiles and the red line represents the median. Whiskers extended to the most extreme data point that is not an outlier.

Experimental Results

Two experiments were undertaken with the aim of establishing whether prioritising agents based on a fitness-proxy could be an effective mechanism for resolving spatial interference. The first experiment uses a population with different speeds, the second experiment adds maximum sensor range as an additional heterogeneous trait which affects how many other agents an ant may consider during a conflict.

Experiment One: Heterogeneous Speeds

In the first experiment, the speed of each agent was unique to create a population that was heterogeneous in one intrinsic trait (speed). The average speed was maintained at 50 for each population size.




Figure 6: $\Delta v = 5$. The food collected by each agent in the population averaged over 50 independent simulations per population size. The heterogeneous trait was speed. *Left*: each ant in the population collected a similar amount of food when slower ants had priority. *Right*: prioritising faster ants produced a greater spread in the amounts collected across a population.

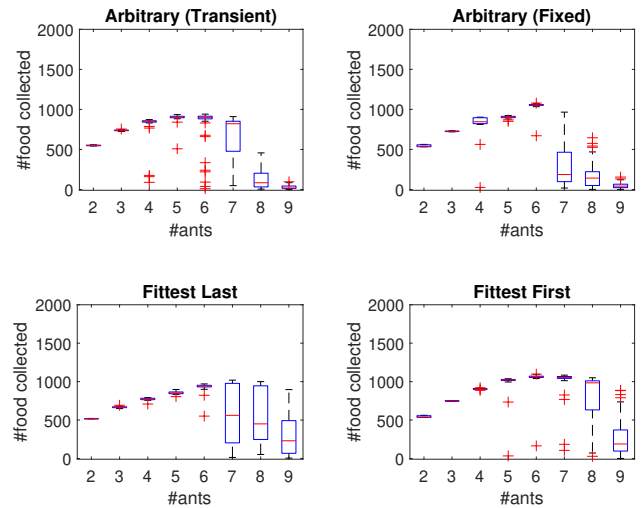


Figure 7: $\Delta v = 10$, The total food collected for different population sizes. The plots show the results from 50 independent simulations per population size. The heterogeneous trait was speed.

Results As the number of agents in the population is increased from 2 to 9, the total amount of food collected initially increases but at a slowing rate, with performance eventually deteriorating (and becoming more variable) beyond a certain population size (fig. 5). It might be expected that prioritising the fastest agent (Fittest First) would produce a population that collects substantially more food but this was not the case. Instead, the performance curves of the four scenarios tested were very similar. The performance, P_i , across the range of population sizes can be summarised as $P_i = \sum_{n=2}^9 F_n^i$ where F_n^i is the total food collected for population size n during simulation replicate $i \in [1, 50]$. Taking the median of P_i for each priority-mechanism shows that Arbitrary Fixed (8807) collected the most food on average over the range of population sizes tested, followed by Fittest Last (8314), Fittest First (8185) and Arbitrary Tran-

sient (7635). A Wilcoxon rank test on the values of P_i revealed some subtleties. Firstly, the Arbitrary Transient results were significant ($p < 0.05$) in pair-wise tests with the other 3 scenarios indicating that a transient hierarchy collects less food on average than a fixed one. Of the fixed hierarchies, Arbitrary Fixed performed better than either Fittest First or Fittest Last ($p < 0.05$), and there was no difference between Fittest First and Fittest Last ($p > 0.05$). Collectively these results suggest a subtle benefit to not always prioritising either the faster or slower agent and to an agent maintaining the same priority over time.

Despite a faster agent having the potential to transport more food in a given time, prioritising speed as a proxy for fitness did not produce the expected gain in performance when resolving interference. This can be seen in the results obtained with 5 ants for Fittest First and Fittest Last where the fastest ant had the potential to move $1.5 \times$ faster than the slowest ant, but giving faster ants priority only resulted in 12% more food being collected by the population.

One reason for this is that using speed as a proxy for fitness results in a near zero-sum scenario. This can be seen by plotting the average food collected by each ant in the population (fig. 6). Prioritising slower agents (Fittest Last) resulted in each ant collecting a similar amount of food. Slower agents were enabled to maximise their individual performance levels, but faster agents were penalised by being forced to move at the speed of slower agents. A greater range of individual performance levels was seen when priority was given to faster agents (Fittest First), but the overall population performance remained roughly similar because improvements for faster agents were almost entirely compensated for by reduced performance levels for slower agents.

Increasing the diversity of the population by doubling the separation between agent speeds within the population decreased performance in all scenarios (fig. 7) which suggests the conflict resolution mechanisms do not scale well to increasing heterogeneity of speed. Calculating the median of the P_i values for each priority mechanism showed Fittest First (6432) now collected more food than the other mechanisms (Wilcoxon rank test on P_i , $p < 0.05$). There was no difference ($p > 0.05$) between Arbitrary Fixed (4470) and Arbitrary Transient (4796), and both these performed worse than Fittest Last (5208). This suggests that there is a benefit to prioritising an agent based on its speed when the range of speeds in the population increases above a threshold, and that benefit is greater when priority is given to the faster agent.

It is also interesting to compare the shape of figures 5 and 7 to the hypothetical curve proposed by Vaughan et al. (2000) and reproduced in fig. 8. Doubling the diversity reduced the height of the peak performance and it occurred at smaller population sizes. Also, the performance did not roll off smoothly as the population size increased. Instead

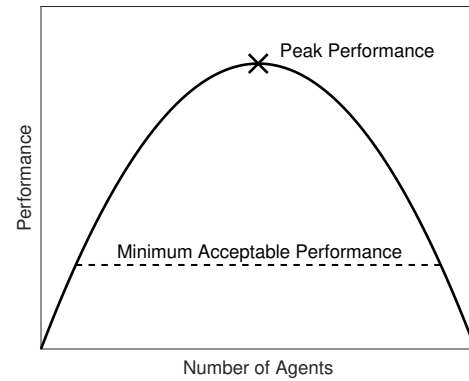


Figure 8: Hypothetical relationship between the size of a population and its performance in the presence of inter-agent interference.

there is a large drop in performance at a population size in the range [7, 9] depending on the priority-mechanism. The curves with lower diversity (fig. 5) suggest a similar shape to the hypothetical curve in figure 8 but the drop in performance did not occur below nine agents.

Experiment Two: Heterogeneous Ranges

Similar to the findings for homogeneous robots reported in Vaughan et al. (2000), the mechanisms for assigning priority gave similar results when speed was used as a proxy for fitness and diversity in the population was low. One possible reason for this is that although agent speed would appear to have a positive effect on task performance, the nature of the task and environment ensured that interference damped the benefit of one agent being faster than another. The significance of a functional trait being dependent on the environment has been seen in the natural world where researchers have compiled lists of which functional traits are significant for particular ecosystems (see Stubbs and Wilson (2004)).

In experiment two, maximum sensor range was chosen as a new trait and proxy for fitness because it is closely coupled to the conflict resolution mechanic. Its effect is to change the maximum number of other agents that an ant can consider during conflict resolution. Ants with a higher maximum have the potential to appreciate the full range of antagonists within larger conflicts. The assumption was that overall system performance would be more sensitive to conflict resolution mechanisms that prioritised this property.

When an ant joins a contest, its sensing range is 10px ensuring that the ant will only be in conflict with its immediate neighbours. This range increases by 5px for every tick the agent remains stationary and decreases by 2px per tick that the agent is moving (down to a minimum of 10px). Therefore, the longer an ant remains stationary, the more ants it can include in its conflict list up to a maximum determined by its assigned maximum sensor range. Figure 4 shows that

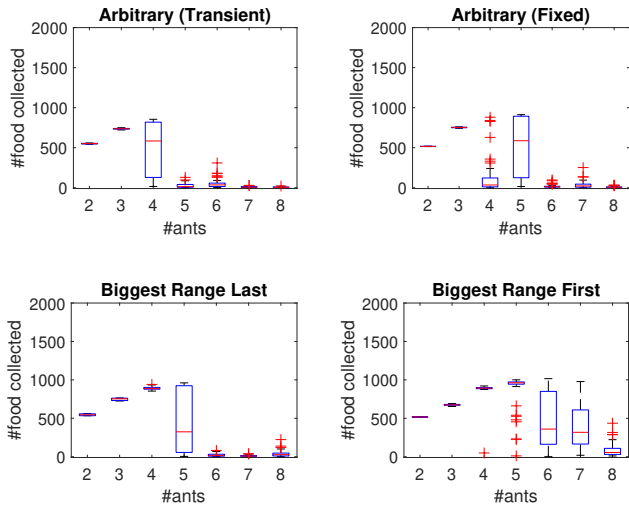


Figure 9: $\Delta v = 10$, The total food collected for different population sizes where priority is assigned based only on maximum sensing range. The plots show the results averaged over 50 repeats of the simulation. Speed is randomised.

if this maximum sensor range is too small then agents may be unable to “see” other antagonists involved in a conflict, resulting in multiple agents becoming Brave and causing a deadlock. In experiment one, the maximum range was the same for all agents (350px) whereas in experiment two each agent was assigned a unique integer value from the range $[20, 10(n + 1)]$ where n is the population size. This created a total of two functional traits in experiment two: speed and maximum sensor range.

Results Figure 9 shows the results of assigning priority based on sensor range as a proxy for fitness. The allocation of a unique speed to each agent is randomised in each simulation. Each of the four scenarios gave a very different set of results, both in terms of the population size for which the most food was collected and the number of ants at which the performance deteriorated. This indicates that performance is more sensitive to sensor range than speed as a fitness-proxy.

Calculating the median of the 50 P_i values for each priority-mechanism and performing a significance test (Wilcoxon rank test on P_i , $p < 0.05$) showed that Biggest First (3910) was the highest performing followed by Biggest Range Last (2617), Arbitrary Transient (2049) and Arbitrary Fixed (1995). The difference between Arbitrary Transient and Arbitrary Fixed was not significant ($p > 0.05$). This indicates that there was a clear advantage to resolving interference using an ordered hierarchy based on maximum sensor range, and there was a further advantage if the hierarchy prioritised agents that could consider more (rather than fewer) agents during a conflict. Interestingly, the significance and ordering of the priority-mechanisms was the same as seen in the $\Delta v = 10$ case for experiment one. This may indicate that a low sensitivity trait (e.g. speed) can display the

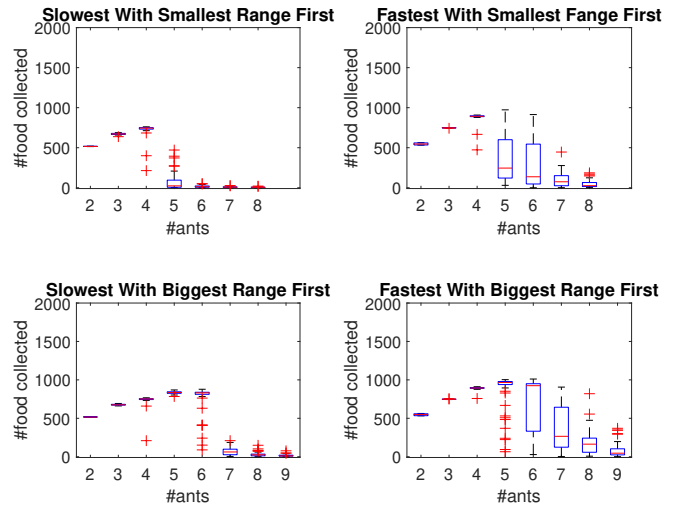


Figure 10: $\Delta v = 10$, The total food collected for different population sizes where priority is assigned based on maximum sensing range and speed. The plots show the results averaged over 50 repeats of the simulation.

same priority-behaviour as a high sensitivity trait when the diversity in the population passes a threshold.

Figure 10 considers scenarios in which speed and maximum sensor range are related together in the population, e.g., agents with longer sensor range are also faster or agents with longer sensor range are also slower.

Again, ordering the scenarios based on the median of the P_i values showed that Fastest With Biggest Range First (4300) collected the most food on average across the range of $[2, 9]$ agents. This was followed by Slowest With Biggest Range First (3686), Fastest With Smallest Range First (3061), and Slowest With Smallest Range First (1988). Each pairing of these results was significant at $p < 0.05$. These results show a positive benefit to prioritising maximum sensor range *and* speed. The results for Slowest with Biggest Range First are the second best performing which adds further evidence to the observation that the mechanism for assigning priority is more sensitive to maximum sensor range than speed.

Summary

In this work, we aimed to test an approach to resolving “same place, same time” spatial interference based on the fitness of an agent. The hypothesis was that assigning priority to an agent that was “fitter” (i.e., higher performing) at the task would produce the best system-level performance. The results demonstrated that this was possible but required an understanding of the environment and inter-agent dynamics. In particular, performance was found to be more sensitive to certain traits, which mirrors how functional traits in natural systems are related to the specific ecosystem that a species inhabits. For the foraging problem presented, the results showed that sensing range (a trait that determined the

maximum number of agents that could be considered within a contest), was more significant as a conflict resolution mediator than the speed of an agent. Assigning priority based on a positive correlation between sensing range and speed was found to amplify the effect, demonstrating that a less significant intrinsic trait can amplify the effect of a significant one. Finally, prioritising the “fitter” individual was found to have a greater relative benefit when the diversity (heterogeneity) in the population increased.

Further Work

This work focused on two traits (speed and sensing range) that are intrinsic to the agent. Previous work has considered either similarly intrinsic properties (Hunt et al., 2019) or more extrinsic properties (Jacyno et al., 2009). Future work should consider how the combination of heterogeneity in both extrinsic and intrinsic agent properties might be exploited within the same multi-agent population.

Increasing the diversity of the population decreased the total amount of food collected for larger population sizes. This raises the issue of relating the way in which a system scales with population *size* to the way that it scales with population *diversity*. What design and analysis tools and techniques could be used to predict and potentially exploit this kind of dynamic variability in a population’s heterogeneity?

Acknowledgements

This work was funded and delivered in partnership between the Thales Group and the University of Bristol, and with the support of the UK Engineering and Physical Sciences Research Council Grant Award EP/R004757/1 entitled ‘Thales-Bristol Partnership in Hybrid Autonomous Systems Engineering (T-B PHASE)’.

This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol (<http://www.bristol.ac.uk/acrc/>).

References

- Balch, T. (2000). Hierarchic social entropy: An information theoretic measure of robot group diversity. *Autonomous Robots*, 8(3):209–237.
- Bennett, C., Bullock, S., and Lawry, J. (2022). A model to investigate exploiting intrinsic multi-agent heterogeneity for spatial interference reduction in an idealised foraging task. <https://doi.org/10.5523/bris.3ota45j8tbx332daufr3vp35ad>.
- Blanckenhorn, W. U. (2000). The evolution of body size: What keeps organisms small? *The Quarterly Review of Biology*, 75(4):385–407.
- 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, S., Zuluaga, M., Zhang, Y., and Vaughan, R. (2005). Rational aggressive behaviour reduces interference in a mobile robot team. In *12th International Conference on Advanced Robotics (ICAR '05)*, pages 741–748. IEEE Press.
- Bullock, S. (2016). “Shit happens”: The spontaneous self-organisation of communal boundary latrines via stigmergy in a null model of the european badger, *Meles meles*. In Froese, T., editor, *Fifteenth International Conference on Artificial Life*, pages 518–525. MIT Press.
- Daly, A. J., Baetens, J. M., and De Baets, B. (2018). Ecological diversity: Measuring the unmeasurable. *Mathematics*, 6(7):119.
- Dugatkin, L. A. (2002). Cooperation in animals: An evolutionary overview. *Biology and Philosophy*, 17(4):459–476.
- Goldberg, D. and Matarić, M. J. (1997). Interference as a tool for designing and evaluating multi-robot controllers. In Kuipers, B. and Webber, B., editors, *14th National Conference on AI (AAAI'97)*, page 637–642. AAAI Press.
- Hooper, D. U., Chapin, F. S., Ewel, J. J., Hector, A., Inchausti, P., Lavorel, S., Lawton, J. H., Lodge, D. M., Loreau, M., Naeem, S., Schmid, B., Setälä, H., Symstad, A. J., Vandermeer, J., and Wardle, D. A. (2005). Effects of biodiversity on ecosystem functioning: A consensus of current knowledge. *Ecological Monographs*, 75(1):3–35.
- Hunt, E. R., Mi, B., Geremew, R., Fernandez, C., Wong, B. M., Pruitt, J. N., and Pinter-Wollman, N. (2019). Resting networks and personality predict attack speed in social spiders. *Behavioral Ecology and Sociobiology*, 73(7):1–12.
- Jacyno, M., Bullock, S., Luck, M., and Payne, T. R. (2009). Emergent service provisioning and demand estimation through self-organizing agent communities. In Sierra, C., Castelfranchi, C., Decker, K. S., and Sichman, J. S., editors, *Eighth International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2009)*, pages 481–488. ACM.
- Liu, W. and Winfield, A. F. T. (2010). Modeling and optimization of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, 29(14):1743–1760.
- Mason, N. W., Mouillot, D., Lee, W. G., and Wilson, J. B. (2005). Functional richness, functional evenness and functional divergence: The primary components of functional diversity. *Oikos*, 111(1):112–118.
- Mayya, S., Pierpaoli, P., and Egerstedt, M. (2019). Voluntary retreat for decentralized interference reduction in robot swarms. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9667–9673. IEEE Press.
- Page, S. (2008). *The Difference*. Princeton University Press.
- Petchey, O. L. and Gaston, K. J. (2006). Functional diversity: Back to basics and looking forward. *Ecology Letters*, 9:741–758.
- Pini, G., Brutschy, A., Birattari, M., and Dorigo, M. (2009). Interference reduction through task partitioning in a robotic swarm. In Filipe, J., Andrade-Cetto, J., and Ferrier, J.-L., editors, *6th International Conference on Informatics in Control, Automation and Robotics*, pages 52–59. INSTICC Press.
- Pitonakova, L., Crowder, R., and Bullock, S. (2018). The information-cost-reward framework for understanding robot swarm foraging. *Swarm Intelligence*, 12(1):71–96.
- Schneider-Fontán, M. and Matarić, M. J. (1998). Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 14(5):815–822.

- Stubbs, W. J. and Wilson, J. B. (2004). Evidence for limiting similarity in a sand dune community. *Journal of Ecology*, 92(4):557–567.
- Tilman, D., Reich, P. B., and Knops, J. M. (2006). Biodiversity and ecosystem stability in a decade-long grassland experiment. *Nature*, 441(7093):629–632.
- Trautman, P. and Krause, A. (2010). Unfreezing the robot: Navigation in dense, interacting crowds. In *Int. Conf. on Intelligent Robots and Systems, IROS 2010*, pages 797–803. IEEE Press.
- Twu, P., Mostofi, Y., and Egerstedt, M. (2014). A measure of heterogeneity in multi-agent systems. In *Proc. 2014 American Control Conference*, pages 3972–3977. IEEE Press.
- Vahl, W. K., Van Der Meer, J., Weissing, F. J., Van Dullemen, D., and Piersma, T. (2005). The mechanisms of interference competition: Two experiments on foraging waders. *Behavioral Ecology*, 16(5):845–855.
- Vaughan, R. T., Støy, K., Sukhatme, G. S., and Mataric, M. J. (2000). Go ahead, make my day: Robot conflict resolution by aggressive competition. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S. W., editors, *Intl. Conf. on Simulation of Adaptive Behavior*, pages 491–500. MIT Press.