Potential of Heterogeneity in Collective Behaviors: A Case Study on Heterogeneous Swarms

Daniela Kengyel¹, Heiko Hamann², Payam Zahadat¹, Gerald Radspieler¹, Franz Wotawa³, and Thomas Schmickl¹

¹ Artificial Life Laboratory at the Department of Zoology, Karl-Franzens University Graz, Graz, Austria, daniela.kengyel@uni-graz.at

² Department of Computer Science, University of Paderborn, Paderborn, Germany, heiko.hamann@uni-paderborn.de

³ Institute for Software Technology, Graz University of Technology, Graz, Austria

Abstract. Research in swarm robotics and collective behaviors is often focused on homogeneous swarms. However, heterogeneity in behaviors can be advantageous as we know, for example, from studies on social insects. Our objective is to study the hypothesis that there are potential advantages of heterogeneous swarms over homogeneous swarms in an aggregation scenario inspired by behaviors of juvenile honeybees. Even without task switching – that is, with predefined, static roles for certain swarm fractions – we find in our case study that heterogeneous swarms can outperform homogeneous swarms for a predetermined set of basic behaviors. We use methods of evolutionary computation to define behaviors imitating those found in honeybees (random walkers, wall followers, goal finders, immobile agents) and also to find well-adapted swarm fractions of different predetermined behaviors. Our results show that non-trivial distributions of behaviors give better aggregation performance.

1 Introduction

In the field of swarm robotics the swarm was, at least early on, defined as (quasi-)homogeneous [3]. This homogeneity referred primarily to the hardware of the swarm robots because that would allow for mass production and consequently for inexpensive swarms. The idea of mass production is followed to date as seen in the kilobot robot [27] and experiments with 1000 robots [28]. Still, a new trend pushes towards heterogeneous swarm research, such as the Swarmanoid project [9]. Heterogeneity in the morphology of swarm members is also seen in natural systems, such as polymorphism in ants [17]. Another concept is to have heterogeneity in the behavior⁴. A swarm with heterogeneous behaviors can still

⁴ The term 'behavior' is used in its biological sense here, that is, it describes a set of an organism's actions. Within the computer science community it would be similar to an agent's strategy or policy and should not be mixed up with low-level behaviors on the level of atomic actions (e.g., turning, gripping).

be homogeneous concerning its hardware, hence, still allowing for mass production. The idea of heterogeneous behaviors in a swarm is that swarm members have predetermined behavioral roles. For example, the polyethism of honeybees [32], that is, taking over different tasks at different times in a bee's life, could be interpreted as a concept of heterogeneous behaviors because the bees rarely switch their general tasks. This is in contrast to frequent task switching in common task allocation and division of labor problems which could be interpreted as temporary heterogeneous behavior [7, 6, 29, 37]. However, each swarm member is typically capable of executing any of the tasks and tasks are switched on demand. There is also the concept of behavioral castes "to describe groups of individuals that perform the same set of tasks in a given period." [10] If tasks are performed for longer periods, agents might specialize, possibly even permanently and hence form heterogeneous behaviors [26]. In natural systems there are examples of swarms that are heterogeneous in their morphology and that have morphology-dependent task switching behaviors [35, 17]. In artificial swarms there are also examples of heterogeneous swarms such the aboved mentioned Swarmanoid swarm robot project [9] or software approaches such as multitype ant colony optimization (MACO) [25]. Swarms that allow for frequent task switching show generally high adaptivity to different work loads and environmental conditions [4,23], but there are also task switching costs (e.g., switching times) that decrease the efficiency [14, 22]. Still, dynamic task switching is advantageous in many situations and, for example, highly developed in many species of social insects. Task partitioning and task switching behaviors are also subject to research in evolutionary swarm robotics [11]. The evolution of heterogeneous behaviors in a multiagent system is reported by [36].

In the following investigations we focus on an extreme case by not allowing any task switching. Agents start with a predetermined behavior and keep it for the whole experiment. The motivation is to simplify the swarm system and to investigate the potential capabilities of such a static non-task-switching system. We hypothesize that swarms with predetermined and fixed behavioral heterogeneity can outperform homogeneous swarms for certain sets of predetermined behaviors. This idea is inspired by the behavior of juvenile honeybees that were found to show several behavioral roles in an aggregation behavior while not switching between them during the whole experiment [34, 20]. Such a complex swarm system with heterogeneous behaviors is an interesting research object in itself but also as an inspiration for how to design swarm robotic systems. We focus on an aggregation task in which the swarm has to find a single target area or to choose between two target areas. In the latter case, the behavior can also be interpreted as a collective-decision making process [34, 15, 12]. This setting is subject to many studies on an algorithm for homogeneous swarms called BEECLUST [30, 5, 16, 13, 2, 1, 18, 19, 31, 21. The BEECLUST algorithm (see Fig. 1) is actually inspired by the above mentioned behavior of young honeybees. Agents controlled by the BEECLUST algorithm move around randomly (step 1 and step 2 create trajectories of straight lines interrupted by rotations due to collision avoidance), whenever they meet another swarm member (step 3) they stop, measure the

- Each agent moves straight until it perceives an obstacle O within sensor range.
- 2.) If *O* is a wall the agent turns away and continues with step 1.
- 3.) If O is another agent, the agent measures the local potential field value. The higher the scalar field value the longer the agent stays still. After this waiting period, the agent turns away from the other agent and continues with step 1.

Fig. 1. The BEECLUST algorithm [30].

local potential field value (e.g., temperature, light, gas concentration), wait for a time proportional to that value, and continue to move randomly afterwards. As a result, the robots form clusters, which is followed by a competition of growing and 'dissolving' robot clusters until one big cluster remains with robots leaving and returning occasionally. The BEECLUST algorithm simplifies the situation found in bees by reducing the different behavior types to only one: random walk. BEECLUST implements a homogeneous approach. The following work can be viewed as an extension of the BEECLUST algorithm to the domain of heterogeneous behavior. In contrast to the study reported in [20], here we investigate behavior compositions with arbitrary numbers that are optimized by evolutionary algorithms, we rely on a mathematical model to represent the individual behavior types, and we investigate different environments.

In this paper, we investigate the above mentioned hypothesis whether a swarm that is heterogeneous in its behavior can outperform a homogeneous swarm under the condition that there are only predetermined basic behaviors and agents are not allowed to switch between them. The motivation is our finding in the behavior of juvenile honeybees that take behavioral roles and never switch them during the run of the experiment [34, 20]. Aggregation at appropriate spots within the bee hive is essential for survival of honeybees and hence we follow that the observed heterogeneous swarm behavior is a well adapted product of natural evolution. In this study, we investigate whether we can reproduce that behavior in simulated agents and test the hypothesis whether heterogeneity outperforms homogeneity in the investigated setting. The results of this study might help to make the right design decisions for systems of swarm robots, such as considering a heterogeneous approach in the first place and then choosing appropriate compositions of predetermined behaviors.

In the following, we limit our case study to a selection of four predetermined behavior types inspired by the biological system of juvenile honeybees. Our study might be considered as an example of biomimicry research due to this choice. However, we also motivate this choice by the opinion that these naturally evolved behaviors might be well adapted to the investigated task of aggregation. The definition of the four behavior types found in juvenile honeybees and a novel model to describe them are our next steps.

2 Four behavior types in juvenile honeybees

Honeybees (Apis mellifera) of age younger than 24 hours show four types of behaviors when allowed to move in a bounded temperature field [34]. The experiments were done in a circular arena surrounded by walls that cannot be climbed by the bees. Heat lamps create a distinct temperature field and it is known that juvenile honeybees have a preference for areas of 36° C [31, 21]. Each of the four behavior types consists of up to two actions: moving and stopping. Except for one type (immobile) all behavior types are combinations of both actions. Switching between the two actions is not considered task switching. The types differ in their movement pattern; there are: random walker (no bias found, neither due to walls nor due to temperature), wall follower (bias towards walls), goal finder (bias towards warmer areas), and immobile agent (no or slow movement only). See Fig. 2 for typical trajectories assigned to their respective behavior type based on tracking data of young honeybees. Note, that the young honeybees never switch between the different behavior types during an experiment.

3 Mathematical model of the behavior types

The behaviors of our agents are directly inspired by the behaviors observed in young honeybees. These behaviors are logically separated in two components: individual behavior aspects differ according to the four identified types and the collective behavior aspects that are identical across all types except for the immobile agents that do not show a reaction to social interactions because they only stay stopped always.

3.1 Individual behavior

We give a general, unified model here that is parametrized to describe all four behavior types. These behavior types are instantiated through different sets of parameters (see Section 3.3). An agent has a position $\mathbf{x} = (x_0, x_1)^{\mathsf{T}}$ (arena limits are $\sqrt{x_0^2 + x_1^2} < 1$), a heading $\phi \in [0, 2\pi)$, and a nominal velocity $v \in [0, 5]$ which is downscaled by discretization to v/100 per time step. An agent can measure an environmental feature, which is temperature in the case of young honeybees but it could also be light, ground color, gas concentration, etc. The environmental feature is modeled by a potential field $P(\mathbf{r})$, $\mathbf{r} \in \mathbb{R}^2$. An agent's turning behavior depends on the environmental feature and/or random effects. The parameter $\alpha \in [0, 1]$ is a weighting factor that determines how intensively an agent follows the gradient of the potential field. A 100% greedy agent following the gradient is defined by $\alpha = 1$. An agent that moves randomly is defined by $\alpha = 0$. Any intermediate value of α defines a corresponding agent that follows the gradient to some extent but is also subject to noise. We define the change of an agent's heading (for simplicity without units) by

$$\frac{d\phi(t)}{dt} = \alpha \min\left(\operatorname{atan}\left(\frac{\partial P(\mathbf{x}(t))}{\partial x_0}, \frac{\partial P(\mathbf{x}(t))}{\partial x_1}\right), \phi_{\max}\right) + (1 - \alpha)\xi(\sigma, t),$$
(1)

for a stochastic process ξ based on Gaussian noise with zero mean, standard deviation σ , and maximal turning angle $\phi_{\text{max}} = 7/18\pi$ ($\phi_{\text{max}} = 70^{\circ}$). An agent's velocity (for simplicity without units) is defined by

$$\frac{d\mathbf{x}}{dt} = \begin{pmatrix} \cos\phi(t)\\ \sin\phi(t) \end{pmatrix} v(t)m(t), \tag{2}$$

for its current nominal speed v(t) and $m(t) \in \{0, 1\}$ giving the agent's current state: m = 0 for *stopped*, m = 1 for *moving*. Note that the nominal speed vis irrelevant in state stopped (m = 0). The transitions between *stopped* and



(a) Tracked bee trajectory (b) Tracked bee trajectory of type *random walker*. of type *wall follower*.



(c) Tracked bee trajectory (d) Tracked bee trajectory of type *goal finder*. of type *immobile agent*.

Fig. 2. Typical tracked trajectories of young honeybees (same length of experiment), assigned to the four behavior types, start of trajectory at triangle, end at circle, 36° C target area at the left hand side of the arena.

moving are modeled as probabilistic state machine with probability to move again P_{move} and probability to stop P_{stop} . Finally, the change of an agent's nominal speed v over time is modeled by a simple Markov chain. The interval of possible velocities [0,5] is discretized as a set of 51 velocities. For each of these discrete velocities v we have a probability of increasing the velocity $P_{\text{incr}}(v)$ by one step (i.e., v' = v + 1/50) and a symmetrical probability of decreasing the speed $1 - P_{\text{incr}}(v)$ (i.e., we force a change). These probabilities P_{incr} define the velocity distribution that results from our model.

3.2 Social behavior

The agents' social behavior, that is the interactions between agents, are homogeneous across all behavior types. They follow the definition of the behavioral model of young honeybees [34] and the definition of the BEECLUST algorithm [30]. Once two agents perceive each other, they stop their motion, measure the local value of the potential field P, and wait for a certain period. This waiting time w is modulated proportionally to the measured potential field value P. It is defined by the function

$$w(P) = \frac{t_{\max}P^2}{\theta + P^2},\tag{3}$$

for parameters $t_{\text{max}} = 132$ time steps and $\theta = 1.4 \times 10^4$. The parameters t_{max} and θ are chosen to generate an appropriate relation between the frequency of robot-robot encounters, the maximum of the potential field P, and the resulting interval of occurring waiting times. During the waiting time all features of the individual behavior are turned off (i.e., velocity v and agent state m are not relevant). Once the waiting time has elapsed the agents do a u-turn of $[-0.25\pi, 0.25\pi]$ and start to move again following their individual behavior type.

3.3 Evolution of parameters for behavior types

Data acquired from experiments with single, young honeybees⁵ are used to find appropriate parameters for our mathematical model. These bee-derived data were manually classified to the four behavior types. A parameter set (σ , α ,

⁵ unpublished, publication in preparation

	RW	WF	GF	IA
σ [radian]	0.090	0.0004	0.57	0.885
α	0.016	0	0.99	0.481
$P_{\rm stop}$	0	0	0.007	0.163
$P_{\rm move}$	0	0	0.024	0.002

Table 1. Typical parameters for the four behavior types of our model: random walker (RW), wall follower (WF), goal finder (GF), and immobile agent (IA).

 $P_{\text{stop}}, P_{\text{move}}, P_{\text{incr}})$ for each behavior type is evolved using a simple genetic algorithm. The population size is 100, we evolve for 100 generations, the mutation rate is 0.25, we select based on proportionate selection, and 30 repetitions per evaluation are done. In each evaluation the agent operates in an arena with only one goal area to avoid side-effects of the symmetrical setting investigated in the swarm experiments. The agent is initially positioned far from that goal area with random orientation and random speed. The agent's behavior is defined by the considered parameter set and it is simulated for 1.5×10^4 time steps. During the simulation all turns and changes of velocity are stored in a histogram of turning angles and a histogram of velocities. The fitness function is a weighted sum of two features: First, it rewards similarities in the histograms of the simulated agent to the histograms acquired from the bee data. Second, type-specific qualities, that are not directly represented by the histograms of turning angles and velocities, are rewarded. In the case of the *qoal finder*, turns towards the goal (i.e., maximum in the potential field P) are rewarded. The gradient of the potential field $\left(\frac{\partial P(\mathbf{x})}{\partial x_1}, \frac{\partial P(\mathbf{x})}{\partial x_2}\right)^{\mathsf{T}}$ defines the optimal direction for each position \mathbf{x} . For each time step, the difference between the agent's direction and the optimal direction is calculated. The sum of these differences is part of the fitness function and hence imposes a minimization problem. In the case of the *wall follower*, time spent close to the wall is rewarded. This is done by defining three areas: a ring-shaped area R_{wall} directly at the wall $\mathbf{x} \in A_{\text{wall}}$: $\sqrt{x_0^2 + x_1^2} > 0.47$, a circular area far from the wall $\mathbf{x} \in \underline{A_{\text{center}}}$: $\sqrt{x_0^2 + x_1^2} < 0.4$, and a second ring in between $\mathbf{x} \in A_{\text{neutral}} : 0.47 < \sqrt{x_0^2 + x_1^2} < 0.4$. In each time step, the agent is rewarded by a score of +1 when positioned on A_{wall} , it receives a penalty of -1 when positioned on A_{center} , and it is treated neutral (± 0) when positioned on A_{neutral} . This score needs to be maximized to evolve a wall following agent. In the case of the *immobile agent*, staying stopped is rewarded which is im-





(b) Histogram of velocities for all 4 behavior types (note logarithmic scale on vertical axis).

Fig. 3. Histograms of turning angles and velocities for all four behavior types based on the mathematical model and parameters as given in table 1 (averaged over 200 repetitions of simulations, 1.5×10^4 time steps each).

plemented by minimizing the agent's average speed. In the case of the *random* walker, no type-specific quality is defined.

The results of these evolutionary runs are shown in table 1 (except for the 50 values of P_{incr}). The resulting histograms of turning angles and velocities (due to $P_{incr}(v)$) for these four behavior types as defined by our model and the parameters given in table 1 are shown in Fig. 3. These results do not allow for a simple interpretation but a few features can be discussed here. The lowest peak for turning angle 0 is found for the random walker which indicates that the turning angle distribution is close to a uniform distribution. The random walker is also one of the fastest. The next peak for angle 0 is that of the goal finder but it also has low values for extreme turning angles. Hence, the goal finder approaches the goal area in a rather straight trajectory. In addition, the goal finder moves slowly. The wall follower has a distribution of turning angles that is close to a uniform distribution similarly to the random walker. However, the maximal turning angle σ is small, which leads to the behavior of a *wall* follower. Additionally there are two more peaks for big turning angles which are the required corrections when following the curved wall around the circular arena. The wall follower moves rather fast. In the case of the *immobile agent* the turning angle is of limited relevance, instead its low average velocity is of more importance.

4 Setup of experiments

In the following experiments, the agents move in a circular arean with either one goal area or two goal areas (see Fig. 4). Following our inspiration from the honeybee experiments we define the potential field P as a temperature field here.



(a) One global goal area at the right hand side of the arena.

(b) A local goal area at the left-hand side and a global goal area at the right-hand side.

Fig. 4. Experimental setup of the arena; the global goal area is located at the righthand side of the arena and contains temperatures between 36° C and 30° C. The local goal area at the left-hand side of the arena contains temperatures between 32° C and 30° C. Each of the goal areas covers 11% of the arena. The potential field P is chosen in a way that there is a global optimum at the right-hand side and an optional local optimum at the left-hand side of the arena. In our experimental settings, these optima are located at the wall, see Fig. 4. In the first experiment 4(a), there is an orange area around the global optimum on the right side. This area is called *global goal area* and has temperatures of 30° C to 36° C ranging from the boundary between the orange and black area to the wall. In the second experimental setup we create a choice-experiment. Additionally to the *global goal area* on the right side, there is an orange area around the local optimum on the left side. This area is called *local goal area* and has a maximal temperature of 32° C at the outer side and 30° C at the boundary between the orange and black area. Each of the goal areas covers 11% of the total arena.

5 Evolution of behavior type compositions

A variation of evolutionary algorithms, called wolf-pack-inspired evolutionary algorithm [38], is used to evolve the composition of behavior types in the swarm. The algorithm maintains overlapping generations and considers a fixed maximum population size. Proportional selection (fitness-based) is used to select individuals (i.e., compositions of behavior types) for mutation that fill empty places in the population. In every generation, one of the individuals, that have not been evaluated yet, is evaluated (alternatively the least evaluated individual if all the individuals have been evaluated already). The algorithm maintains the hierarchy in the population and keeps its diversity by removing older individuals with an equal or lower fitness than a newly evaluated individual (with a probability factor). The fitness function is defined by

$$F = G - L \tag{4}$$

where G is the number of agents within the global goal area and L is the number of agents within the local goal area (if there is one).

6 Results

For both experimental settings (one global goal and choice-experiment), we investigate the potential of heterogeneous swarms. As described in the above section, we use evolutionary algorithms to adapt the swarm's behavior-type composition to the environment. The experiments are based on a fixed swarm size N = 15. The results are based on n = 18 independent runs of the evolutionary algorithm and the population of compositions was initialized to a random uniformly distributed setting of behavior types. The evolved approach is compared to the fitness of several homogeneous swarm settings (Fig. 5) that were evaluated in n = 100 independent simulation runs (no evolution because composition is predetermined). In the first three homogeneous swarm settings we use a swarm size of N = 15. For the last two settings we used a swarm size of N = 12 to test for a potential density dependency.

First we focus on the experiment with only one goal area (Fig. 5(a)). The median fitness for 15 random walkers is 7, for 15 goal finders it is 5, for 15 wall followers it is 9, for 12 random walkers it is 6, and for 12 wall followers it is 8. For the heterogeneous swarm optimized by evolution the median fitness is 10 (n = 18). The evolved behavior-type composition is found to be significantly better than the homogeneous swarms (based on Wilcoxon rank sum test, p < 0.05).

Figure 5(b) shows the results of the choice experiment (global goal area and local goal area). Here we compare the evolved heterogeneous behavior-type composition (first box plot, labeled 'Evo') to homogeneous behavior-type compositions. The median fitness for 15 random walkers is 2, for 15 goal finders it is 0, for 15 wall followers it is 3, for 12 random walkers it is 2, and for 12 wall followers it is 2. For the heterogeneous swarm optimized by evolution the median fitness is 5.5 (n = 18). The evolved behavior-type composition is found to be significantly better than the homogeneous swarms (based on Wilcoxon rank sum test, p < 0.05). Hence, our heterogeneous approach is the most effective variant of all tested configurations. The results for 12 random walkers and 12 wall followers indicate no dependency on density. The motivation of this test is based on results we report below and the consideration that immobile agents might potentially be used to virtually decrease the agent density.

Next, we evolve behavior-type compositions for different environments (one or two goal areas) and different initializations of the composition populations. We start with the setting that has only one goal area (see Fig. 4(a)). The evolution-



(a) Setting with one goal area at the wall. (b) Setting with a *global goal area* and a *local goal area* (choice experiment).

Fig. 5. Comparison of the best fitness between one evolved heterogeneous setting and several homogeneous swarm settings for one goal area at the wall (left) and the choice experiment on the right (global goal area and local goal area); heterogeneous swarm (labeled 'Evo'), homogeneous swarms with only random walkers (RW), goal finders (GF), or wall followers (WF). In both settings, the heterogeneous swarm is significantly better than all homogeneous swarms (based on Wilcoxon rank sum test, p < 0.05). Other significances are not shown.

ary approach is as described above, that is, the initial population of compositions is sampled from a random uniform distribution. For our analysis, we take the best composition of the last population from each evolutionary run. The box plots shown in Fig. 6(a) give a summary of these best compositions. The number of occurrences for each behavior type is given for the n = 18 best evolved compositions. The median number of goal finders is 1.5, the median of wall followers is 8.5, the median of random walkers is 2.5, and the median of immobile agents is 1. It is counterintuitive that goal finders are relatively infrequent while the high number of *wall followers* might seem reasonable because the goal area is located at the wall. In Fig. 6(b) we give an overview of the type frequencies of the current best compositions over the number of evaluations averaged over all evolutionary runs. We started with compositions that are in average uniformly distributed. During the first 10 evaluations the number of *immobile agents* is decreased while the number of *wall followers* is increased quickly. The number of random walkers increases initially but then decreases again. The number of goal finders is decreased over a long period during the first 40 evaluations. After about 100 evaluations a saturation effect is observed.

Next we investigate the choice experiment (local goal area on the left side and a global goal area on the right side of the arena). The box plots of Fig. 7(a) give the number of agents for each behavior type as they occurred in the best compositions of n = 20 independent evolutionary runs. The median number of goal finders is 1, the median of wall followers is 6, the median of random walkers is 2.5, and the median of immobile agents is 4. As expected the number of goal finders is smaller in comparison to the setting with only one goal (cf. Fig. 6(a)) because goal finders merely follow the gradient and the swarm separates between





(a) Results of the evolution with one goal area. The plot shows the median amount of behavioral types that are used to compose a heterogeneous swarm with the highest fitness.

(b) Type frequencies of the best composition over evaluations

Fig. 6. Results of evolved swarm compositions with one goal area and all four behavioral types: random walker (RW), goal finder (GF), wall follower (WF), and immobile agent (IA).

the two goal areas. The number of *wall followers* is decreased, the number of *random walkers* is increased in its variance, and the number of *immobile agents* is increased in comparison to the one-goal setting. Especially the increase of *immobile agents* is counterintuitive because they are of no direct use to maximize the fitness function. In Fig. 7(b) we give an overview of the type frequencies of the current best compositions over the number of evaluations averaged over all evolutionary runs. Starting from approximately uniformly distributed compositions the number of *immobile agents* first decreases and is then increased slowly over about 120 evaluations at the cost of random walkers. After about 130 evaluations a saturation effect is observed.

7 Discussion

Concerning the results for the one-goal setting (Fig. 5(a) and 6) one would expect that the best fitness in this setup is achieved by making exclusive use of *goal finders* only. From our experience with the simulation we can tell that too many *goal finders* actually block each other in areas before the goal area which results in clusters outside of the goal area. Instead, a limited number of *goal finders* turns out to be useful because such deadlock situations are then avoided. They serve as seeds within the goal area and help agents of other types to form clusters inside the goal area more easily, which is an example of how the different behavior types create opportunities of cooperation between agents. Most of the agents of the evolved heterogeneous swarms are *wall followers*. With only one goal area present, the *wall followers* always end up in the goal area and form a cluster. In comparison, the number of *random walkers* is low. Their approach to the goal area is slower because they might form clusters within the center of the arena. Eventually, they join the cluster in the goal area and join the *wall followers*. Therefore, in this setup a high amount of *wall followers* is the



(a) Median number of behavior types as they occur in the best swarm composition.

(b) Development of the best composition that is evaluated at each time step.

Fig. 7. Results of the evolution for the choice experiment (*local goal area* and *global goal area*) and all four behavioral types: random walker (RW), goal finder (GF), wall follower (WF), immobile agent) (IA).

better choice. In an extended study, that is in preparation, we have also done experiments with goals not positioned at the walls. The number of *wall followers* decreases for that setting as expected but the qualitative result of our study is not influenced by the positions of the goal areas.

Concerning the results for the choice experiment (Fig. 5(b) and 7) the small number of *qoal finders* is explained by the fact that they are not able to distinguish between a global and a local goal area because they merely follow the local gradient. Hence, they are not able to increase fitness (F = G - L). This is also indicated by the zero median for homogeneous *goal finder* swarms (Fig. 5(b)). Still, *qoal finders* might be useful in a heterogeneous swarm to mark the goal areas and to serve as social seeds that attract others. Compared to the results of the experiment with only one goal area, the median amount of *immobile agents* is higher. Intuitively it seems inappropriate to use any *immobile agent* because they never enter the goal area when placed outside of it initially. However, they are part of many evolved swarm compositions although the optimization algorithm is effective [38] and we also do not enforce that all four behavior types have to be included in the solution. Thus, additional experiments are required to investigate the role of *immobile agents* and to find a sound explanation of why *immobile agents* are useful for the swarm in both our model and also in the natural swarms of honeybees. We can only speculate that *immobile agents* might have the functionality of a barrier and might slow down or even block agents that switch between goals. That way *immobile agents* might prevent other agents from visiting the *local goal area* and hence might stabilize the whole decisionmaking process. However, this requires more investigations and will be done in future work.

8 Conclusion

In this paper we have investigated swarms of agents that are heterogeneous in their behaviors. The idea is to simplify the swarm system by predefining static roles for certain swarm fractions. Even for the investigated extreme case without task switching, the heterogeneous swarm outperforms homogeneous swarms in the investigated aggregation scenario for the selected, predetermined behavior types. For now, all our results are based on one set of predetermined behaviors and one kind of collective task. However, the selection of behaviors was not arbitrary but inspired by results from biological experiments with juvenile bees. Still, the generalization of this work is left for future work.

The evolved compositions of behavior types indicate a rather complex underlying system that creates nontrivial distributions of behaviors which might even be perceived as counterintuitive. While the behavior types themselves were simple and predefined here, it is of course an option to determine the behavior types also by evolutionary computation or other methods of machine learning. However, for applications of swarm robotics, such as nanorobotics [24], it is attractive to make use of simple predetermined behaviors.

The effectivity of the evolved behavior compositions is certainly interesting, raises questions, and allows for different interpretations. While the four behavior types all score low in homogeneous swarms, they allow for a much more efficient aggregation behavior once combined. Obviously cooperation among different types is crucial and teamwork of a diverse team is essential. A tempting interpretation is that the results might be compared to findings in natural swarms that rely on certain degrees of leadership [8]. Only leadership is difficult to define here. The goal-oriented and greedy behavior of the *goal finder* is not helpful for the swarm per se. It requires a *random walker* and a *wall follower* to make use of the social seed within the goal area created by a *goal finder*. Hence, we observe a sophisticated interplay of agents with different approaches and capabilities that outperform their homogeneous counterparts as a heterogeneous swarm.

Also note that the use of simulations is potentially the only means to investigate the concept of predetermined behavioral roles in the natural complex system of young honeybees. Following the common standards of experiment design in biology it is not an option to use the same subjects (bees) in several replications of the experiment. In our case here an initial experiment would be necessary to label the bee with its behavioral role and in a second experiment we could create the desired swarm composition of behavior types. However, the bee might be influenced by the initial experiment and show a different behavior. Hence, simulations are a useful tool to investigate this complex system of interacting honeybees.

The results of this study support a core idea of swarm robotics that the interplay of several simple behaviors generates complex behaviors due to multiple interactions. This case study's main result is that heterogeneous swarms based on predetermined behaviors without task switching can perform well. Our approach is not limited to the study of the BEECLUST algorithm. Also other collective behaviors can be explored, such as heterogeneity in the stimulus-response functions of bees in their waggle-dance behavior [33]. In future work, we plan to do a complete sensitivity analysis of the many paramters in our model. In addition, we plan to work our way towards a generalization of our approach, for example, allowing different sets of predetermined or even learned behaviors. Although this study was guided by the biological inspiration of young honeybees' behavior, our future research will focus more on engineering applications of heterogeneous swarms in (evolutionary) swarm robotics.

9 Acknowledgements

This work is supported by: Austrian Science Fund (FWF): P19478-B16 and P23943-N13 (REBODIMENT); EU FP7 FET-Proactive 'ASSISI_{bf}', no. 601074; EU H2020 FET 'flora robotica', no. 640959.

References

1. Arvin, F., Turgut, A.E., Bazyari, F., Arikan, K.B., Bellotto, N., Yue, S.: Cue-based aggregation with a mobile robot swarm: a novel fuzzy-based method. Adaptive

Behavior 22(3), 189-206 (2014)

- Arvin, F., Turgut, A.E., Yue, S.: Fuzzy-based aggregation with a mobile robot swarm. In: Swarm Intelligence (ANTS'12). Lecture Notes in Computer Science, vol. 7461, pp. 346–347. Springer, Berlin (2012)
- Beni, G.: From swarm intelligence to swarm robotics. In: Şahin, E., Spears, W.M. (eds.) Swarm Robotics - SAB 2004 International Workshop. LNCS, vol. 3342, pp. 1–9. Springer-Verlag, Santa Monica, CA (Jul 2005)
- Berman, S. et al.: Optimized stochastic policies for task allocation in swarms of robots. Robotics, IEEE Transactions on 25(4), 927–937 (2009)
- Bodi, M., Thenius, R., Szopek, M., Schmickl, T., Crailsheim, K.: Interaction of robot swarms using the honeybee-inspired control algorithm beeclust. Mathematical and Computer Modelling of Dynamical Systems 18(1), 87–100 (2012)
- Bonabeau, E., Dorigo, M., Theraulaz, G.: Swarm Intelligence: From Natural to Artificial Systems. Oxford Univ. Press (1999)
- Campbell, A., Wu, A.S.: Multi-agent role allocation: issues, approaches, and multiple perspectives. Autonomous Agents & Multi-Agent Systems 22, 317–355 (2011)
- Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A.: Effective leadership and decision-making in animal groups on the move. Nature 433, 513–516 (Feb. 2005)
- 9. Dorigo, M., et al.: Swarmanoid: a novel concept for the study of heterogeneous robotic swarms. IEEE Robotics & Automation Magazine 20(4), 60–71 (2013)
- Dorigo, M., Bonabeau, E., Theraulaz, G.: Ant algorithms and stigmergy. Future Generation Computer Systems 16(9), 851–871 (Jun 2000)
- Ferrante, E., Dúeñez Guzman, E., Turgut, A.E., Wenseleers, T.: Evolution of task partitioning in swarm robotics. In: et al., V.T. (ed.) Proceedings of the Workshop on Collective Behaviors and Social Dynamics of the European Conference on Artificial Life (ECAL 2013) (2013)
- Garnier, S., Gautrais, J., Asadpour, M., Jost, C., Theraulaz, G.: Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. Adaptive Behavior 17(2), 109–133 (2009)
- 13. Hamann, H.: Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. Swarm Intelligence 7(2-3), 145–172 (2013)
- Hamann, H., Karsai, I., Schmickl, T.: Time delay implies cost on task switching: A model to investigate the efficiency of task partitioning. Bulletin of Mathematical Biology 75(7), 1181–1206 (July 2013)
- Hamann, H., Meyer, B., Schmickl, T., Crailsheim, K.: A model of symmetry breaking in collective decision-making. In: Doncieux, S., Girard, B., Guillot, A., Hallam, J., Meyer, J., Mouret, J. (eds.) From Animals to Animats 11. LNAI, vol. 6226, pp. 639–648. Springer-Verlag (2010)
- Hereford, J.M.: Analysis of BEECLUST swarm algorithm. In: Proc. of the IEEE Symposium on Swarm Intelligence (SIS 2011). pp. 192–198. IEEE (2011)
- 17. Hölldobler, B., Wilson, E.: The ants. Belknap Press of Harvard University (1990)
- Kengyel, D., Schmickl, T., Hamann, H., Thenius, R., Crailsheim, K.: Embodiment of honeybee's thermotaxis in a mobile robot swarm. In: 10th European Conference on Artificial Life (ECAL'09). LNCS, vol. 5777/5778. Springer-Verlag (2011)
- Kengyel, D., Thenius, R., Crailsheim, K., Schmickl, T.: Influence of a social gradient on a swarm of agents controlled by the beeclust algorithm. Advances in Artificial Life, Proceedings of the 12th European Conference on the Synthesis and Simulation of Living Systems, ECAL13 12, 1041–1048 (2013)
- Kengyel, D., Wotawa, F., Schmickl, T.: Towards swarm level programming: The role of different movement patterns in swarm systems. Swarm Intelligence (2014), submitted

- Kernbach, S., Thenius, R., Kornienko, O., Schmickl, T.: Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic swarm. Adaptive Behavior 17, 237–259 (2009)
- 22. Khaluf, Y., Birattari, M., Hamann, H.: A swarm robotics approach to task allocation under soft deadlines and negligible switching costs. In: del Pobil, A.P., Chinellato, E., Martinez-Martin, E., Hallam, J., Cervera, E., Morales, A. (eds.) Simulation of Adaptive Behavior (SAB 2014). LNCS, vol. 8575, pp. 270–279. Springer (2014)
- Labella, T.H., Dorigo, M., Deneubourg, J.L.: Division of labor in a group of robots inspired by ants' foraging behavior. ACM Transactions on Autonomous and Adaptive Systems (TAAS) 1(1), 4–25 (2006)
- 24. Lenaghan, S., Wang, Y., Xi, N., Fukuda, T., Tarn, T., Hamel, W., Zhang, M.: Grand challenges in bioengineered nanorobotics for cancer therapy. IEEE Transactions on Biomedical Engineering 60(3), 667–673 (2013)
- Liu, X., Li, X., Shi, X., Huang, K., Liu, Y.: A multi-type ant colony optimization (maco) method for optimal land use allocation in large areas. International Journal of Geographical Information Science 26(7), 1325–1343 (2012)
- Lorenz, K.: Vergleichende Verhaltensforschung: Grundlagen der Ethologie. Springer (1978)
- Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., Nagpal, R.: Kilobot: A low cost robot with scalable operations designed for collective behaviors. Robotics and Autonomous Systems 62(7), 966–975 (2014)
- Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. Science 345(6198), 795–799 (2014),
- Schmickl, T., Crailsheim, K.: TaskSelSim: a model of the self-organization of the division of labour in honeybees. Mathematical and Computer Modelling of Dynamical Systems 14, 101–125 (2008)
- Schmickl, T., Hamann, H.: BEECLUST: A swarm algorithm derived from honeybees. In: Xiao, Y. (ed.) Bio-inspired Computing and Communication Networks. CRC Press (March 2011)
- Schmickl, T., Thenius, R., Möslinger, C., Radspieler, G., Kernbach, S., Crailsheim, K.: Get in touch: Cooperative decision making based on robot-to-robot collisions. Autonomous Agents and Multi-Agent Systems 18(1), 133–155 (February 2008)
- 32. Seeley, T.D.: Adaptive significance of the age polyethism schedule in honeybee colonies. Behavioral Ecology and Sociobiology 11, 287–293 (1982)
- Seeley, T.D.: Honey bee foragers as sensory units of their colonies. Behavioral Ecology and Sociobiology 34, 51–62 (1994)
- Szopek, M., Schmickl, T., Thenius, R., Radspieler, G., Crailsheim, K.: Dynamics of collective decision making of honeybees in complex temperature fields. PLoS ONE 8(10), e76250 (10 2013)
- Wilson, E.: The relation between caste ratios and division of labour in the ant genus *Pheidole* (Hymenoptera: Formicidae). Behav. Ecol. Sociobiol. 16, 89–98 (1984)
- 36. Yong, C.H., et al.: Coevolution of role-based cooperation in multiagent systems. IEEE Transactions on Autonomous Mental Development 1(3), 170–186 (2009)
- Zahadat, P., Crailsheim, K., Schmickl, T.: Social inhibition manages division of labour in artificial swarm systems. In: Lio, P., Miglino, O., Nicosia, G., Nolfi, S., Pavone, M. (eds.) 12th European Conference on Artificial Life (ECAL 2013). pp. 609–616. MIT Press (2013)
- 38. Zahadat, P., Schmickl, T.: Wolfpack-inspired evolutionary algorithm and a reaction-diffusion-based controller are used for pattern formation. In: Proceedings of the 2014 Conference on Genetic and Evolutionary Computation. pp. 241–248. GECCO '14, ACM, New York, NY, USA (2014)