

Collective Decision Making in a Swarm of Robots: How Robust the BEECLUST Algorithm Performs in Various Conditions

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

Thomas Kunzfeld
Artificial Life Laboratory at the
Department of Zoology
Karl-Franzens University Graz,
Austria
thomas.kunzfeld@uni-
graz.at

Payam Zahadat
Artificial Life Laboratory at the
Department of Zoology
Karl-Franzens University Graz,
Austria
payam.zahadat@uni-
graz.at

Thomas Schmickl
Artificial Life Laboratory at the
Department of Zoology
Karl-Franzens University Graz,
Austria
thomas.schmickl@uni-
graz.at

ABSTRACT

In this paper a honeybee inspired collective-decision-making algorithm called BEECLUST is studied in a swarm of autonomous robots and the performance of the swarm is investigated in different conditions. The algorithm has low requirements thus it is promising for implementation in robots with low resources. Here the algorithm is applied in swarms of improved e-puck robots in three different conditions in order to study the strengths and limitations of the algorithm. The collective system demonstrated a high performance in adapting to a dynamic environment as well as a very low sensitivity to additional robots with malfunctioning sensors. On the other hand the system shows a strong response to robots that act as social seeds influencing the decision-making of the swarm.

Categories and Subject Descriptors

F.2.m [Analysis of Algorithms and Problem Complexity]: Miscellaneous; I.2.9 [Artificial Intelligence]: Robotics—Autonomous vehicles; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Multiagent systems

General Terms

Algorithms, Experimentation, Performance

Keywords

swarm robotics, autonomous agents, self-organisation, bio-inspired algorithm, swarm intelligence, collective decision making, search

1. INTRODUCTION

Social insects are promising sources of inspiration for collective intelligence and swarm robotics where simple rules implemented by the agents lead to complex behaviour of the collective system. In this paper an algorithm inspired by young honeybees is implemented in real embodied robots. The algorithm is investigated for its behaviour in different conditions in order to make a deeper understanding of the capabilities and limitations of the algorithm for swarm robotics and getting insights into mechanisms used in natural honeybees.

Young honeybees are not capable of flying. They can move around in the hive and they are capable of collectively finding locations with the best temperature [5, 18]. It has been shown [7, 9], that there is a preferred temperature (36°C) that is an important factor for the development of the young honeybees. This is the main temperature in the brood nest of a bee-hive and acts as a mechanism to confine the brood in the brood nest [16].

Previous experiments demonstrate that individual young honeybees are mostly not capable of locating the area with the preferred temperature but swarms of them can find the preferred spot collectively [15, 17, 19]. This collective behaviour is modelled by the BEECLUST algorithm [15] which consists of simple rules for the individual agents and leads to the collective behaviour of the swarm finding the global optimum out of several local optima in the environment. The algorithm does not require any explicit communication between agents or directional sensing of the environmental modalities such as temperature. It works based on local measurements and the capability of agents to distinguish be-

tween obstacles and other agents in the near vicinity. The minimal requirements of the algorithm make it promising for applications in swarm robotics where several robots with limited capabilities are supposed to perform a task collectively. The BEECLUST algorithm has been analysed [11, 10, 8], investigated for improvements [2, 4, 3, 12], and applied [17, 15, 13] in previous works in simulation and also with robots in a light gradient. In this paper we implement the algorithm in swarms of autonomous e-puck robots [6]. Since the BEECLUST algorithm is derived from the behaviour of young honeybees, experimenting with a temperature gradient gets us closer to the situation that young honeybees are faced with, thus might be interesting also from a biological point of view. Therefore, we extended the e-puck robots with temperature sensors. A set of experiments were performed in an arena containing two heat sources as global and local optima. The behaviour of the collective system is investigated in a set of collective choice experiments with different conditions. In the first scenario, the behaviour of the system is investigated in a dynamic environment where the global optima changes its location to the local optimum to show that the agents are able to react on environmental changes under sophisticated conditions. The second and third scenario deal with the situations where some of the robots do not act as normal robots. In the second scenario the effect of social seeds is investigated. Social seeds are robots that try to influence the swarm decision by locating themselves in a certain area (i.e. local optimum) of the arena. Here we want to investigate how strong the effect of a social seed is. The third scenario investigates whether or not robots with broken sensors can reduce the performance of the swarm of intact robots and thus, harm the robots decision making process. This is done by adding robots with broken sensors to a swarm of fully functional robots.

2. MATERIAL & METHOD

2.1 Hardware platform

As a hardware platform, the “e-puck” robot is used [6]. The e-puck features a broad array of different sensors, however, only the infrared (IR) proximity sensors are used here for obstacle detection and detection of another robot. As the e-puck has no temperature sensors, we designed an extension board which provides additional sensors like color sensors, light sensors and temperature sensors, whereas only the temperature sensors are used here. Temperature measurement is performed by a MAXIM MAX6636 integrated circuit which supports up to 6 external temperature sensors. For our experiments we only use a single temperature sensor which measures the temperature slightly above the ground. Figure 2.1 shows an e-puck robot with its extension board and temperature sensor.

2.2 BEECLUST algorithm

The BEECLUST algorithm is a swarm algorithm inspired by the behaviour of young honeybees. The algorithm enables the swarm to find the global optimum out of several local optima. The algorithm is based on collisions between the agents and the measurement of local modality by individual agents. The algorithm is summarised as follows:

- Each agent moves randomly until a collision is detected.

- If the collision has happened with an obstacle, the agent turns around and continues with the random movement.
- If the collision has happened with another agent, the agent measures the local environmental modality (e.g., temperature) and waits in its position for a certain period of time. The waiting time is a monotonic function of the measured value of the local modality.

2.3 Setup

In the following experiments, a typical setup for a choice experiment is used. A circular arena (figure 2(a)) of $r = 2m$ is set up with two heat lamps located in two sides of the arena (semi-circles) making a global optimum and a local optimum temperature. A picture of the arena taken with an infrared camera is shown in figure 2(b). On the right side, there is the *global optimum area* (green with yellow area in the middle). The *global optimum* has a maximum temperature of $36^{\circ}C$ at the middle and $30^{\circ}C$ at the boundaries. The *local optimum area* is located on the left side of the arena (green area at the left side). The *local optimum* has a maximum temperature of $32^{\circ}C$ at the middle and $30^{\circ}C$ at the boundaries. The two areas cover 11% of the arena, each. The ambient temperature of the room, thus the temperature between the *global optimum* and the *local optimum*, is approx. $28^{\circ}C$. At the beginning of each experiment, the robots are released with a random heading of $[0\pi, 2\pi]$ in the middle of the arena (area with about $28^{\circ}C$).

2.4 Experiments

In the following sections we describe the three sets of experiments with the swarms of e-puck robots running the BEECLUST algorithm as the controller of the robots. The three different scenarios are as follows: dynamic environment, social seed scenario, and impaired sensor robots. As a control experiment we conducted the classical binary choice setup with 10 robots for the dynamic environment and social seed scenarios. The control experiment for the impaired sensor scenario is the classical binary choice setup with only three robots. Every scenario is repeated $n = 10$ times. In each experiment, where the number of robots in certain areas are investigated, the robots are counted every 20 seconds and the median of the counted values are taken as the number of robots in that area for that experiment.

2.4.1 Dynamic environment

In [17] it was shown in a light gradient, that agents controlled by the BEECLUST algorithm are not only able to find the global optimum out of several local optima, but they are also able to react on environmental changes like changing the global and local optima. Here we want to show that it is also possible under more sophisticated environmental conditions like in a temperature gradient.

The experiment is carried out in two steps. The setup in the first step is identical to what is described in Sec. 2.3 with a *global* and *local optima* located respectively at the right and left side of the arena. The first step consists of a single observation phase and lasts for 15 minutes. At the beginning of the first step, agents are released in the middle of the arena. After 15 minutes, the second step starts with switching off the heat source of the *global optimum*. The second step consists of two phases which last for 5 and 10




Figure 1: e-puck with the extension board (on top) and a temperature sensor (at the left side slightly above ground).

minutes respectively. In the first 5 minutes the *global optimum* cools down to the ambient temperature (28°C). Then in the next 10 minutes the formerly called *global optimum* keeps a median temperature of 28°C and the formerly called *local optimum* becomes the new global optimum. Without changing the heat source of the former local optimum, its temperature decreases to 30°C due to the absence of the heat source of the former global optimum. The duration of a complete experiment is 30 minutes. See figure 4 for the changes in the median temperature during the three phases.

2.4.2 Social seed

It has been previously shown in [1] that adding a social component to individual learning can improve the performance of a population of autonomous robots developing a simple group behaviour, i.e. foraging. Here we used a social seed as the social component in the arena. The social seed is a robot that is immobile but its social effect is the same as other robots meaning that other robots can perceive its presence and react to it as a robot. In a previous work [14] we investigated the effect of social seeds in simulation. Here we are interested in investigating if this effect also exists in robots controlled by the BEECLUST algorithm and how strong the effect is compared to the results from the simulation experiment. Therefore, we created a setup similar to the simulation setup: The main swarm consists of 10 robots. In the *local optimum* we placed an immobile robot as a social seed which can be recognised by other robots (see figure 3). To keep the conditions fair for the two optima (e.g. in terms of available space), we placed a dummy robot in the *global optimum*. The dummy robot is perceived as an obstacle by other robots.

2.4.3 Impaired sensors

During experiments with swarm robots, we raised the question if robots that are not fully functional do harm the performance of the swarm consisting of fully functional robots or if they can still help other swarm members. This is especially relevant if the presence of robots in the arena has

other purposes in addition to finding the *global optimum*.

In the experiment here, the performance of a swarm of intact robots aiming to find the *global optimum* is investigated in presence of additional robots with malfunctioning temperature sensors. Here, the malfunctioning temperature sensors measure a temperature of 0°C and thus, have always a waiting time of $t = 0\text{s}$. All the robots are controlled by the BEECLUST algorithm. The impaired robots cannot measure the temperature and therefore never stop, but they can still create collisions and trigger other robots to measure the local temperature.

The intact swarm in this experiment consists of 3 fully functional robots. To measure the effect of impaired robots, we added 7 robots that are not able to measure the temperature. In the control experiment, the 3 fully functional robots were used. We used only 3 robots here, because the aggregation with 3 robots is stable enough to measure a decrease of performance if the impaired robots harm the decision making process. On the other hand, to measure if the performance can be increased with the help of impaired robots, the aggregation of 3 robots is instable enough so that there is potential to improve the aggregation.


3. RESULTS

3.1 Results: Dynamic environment

The dynamic environment experiment was created to demonstrate that agents controlled by the BEECLUST algorithm are flexible in their decision-making in a dynamic environment even where the environment is very inert and unstable (like the temperature-gradient). Figure 4 shows the changes in temperature over time. In phase 1, both optima, *global optimum area* and *local optimum area* are present. After 15 minutes, the heat source of the *global optimum area* is switched off and a cooling phase of 5 minutes starts (phase 2). Then only one optimum is present (phase 3). In figure 5 the corresponding distribution of agents over time is shown. In phase 1 all robots start in the middle of the arena. Af-



(a) Setup of the arena. Blue lines define the *local/global optimum areas*.



(b) Heatmap of the arena with two sources of heat. On the left side the *local optimum area* with 32°C and on the right side the *global optimum area* with a maximum of 36°C.

Figure 2: Experimental Setup of the arena and the corresponding heatmap.



Figure 3: The figure shows the setup of the social seed. Left: *local optimum* area with a social seed (immobilised robot). Right: *global optimum* area with a dummy robot to ensure similar circumstances of the two optima areas.

ter a short period, most of the agents are aggregated in the *global optimum*. In phase 2, the *global optimum* is shut down and therefore the cluster dissolves. After a few minutes, the robots start to cluster in the *local optimum* which is now the new global optimum in the arena. At the end of the experiment most robots are clustered in the formerly called *local optimum*.

Figure 6 shows the median number of robots in the *global optimum* and *local optimum* for the different phases. In phase 1 the median number of robots in the *global optimum* is 7 and in the *local optimum* is 0. The median number of robots in the cooling phase is 6.5 and 0 in the *global optimum* and in the *local optimum*, respectively. In the last 10 minutes of the experiment (phase 3) the median number of robots in the *global optimum* is 0 and in the *local optimum* is 6. Statistical significances are tested with the Wilcoxon-Test (*global optimum* against *local optimum*) and the U-Test (number of robots in the optimum area of one phase against the same optimum area of another phase) with $p < 0.05$. All boxplots in the figure are significantly different to the other boxplots, except *global optimum* of phase 1 with *global optimum* of phase 2, and also *local optimum* of phase 1 with *local optimum* of phase 2.

3.2 Results: Social seed

In figure 7 the median time that each agent spends in the *global optimum* area and in the *local optimum* area are depicted for both control experiment (no social seed) and the experiment with social seed. In the control experiment, the robots spend 67% of the time in the *global optimum* area and 4.6% in the *local optimum* area (left side of figure 7). In the experiment with social seed (right side of figure 7), the median time each agent spends in the *global optimum* area is 26% and in the *local optimum* area is 45.8%. Statistical significances are tested with the Wilcoxon-Test (*global optimum* tested against *local optimum*) and the U-Test (*global optimum* of one experiment against *global optimum* of another experiment) with $p < 0.05$. Note that here we compare the median time the robots spent in the respective optimum area in % to allow comparison of our results to simulation results reported in [14].

3.3 Results: Impaired sensors




Figure 4: Course of temperature in the different phases over time in the dynamic environment experiment.




Figure 5: Course of experiment in the dynamic environment: Median number of robots in the different areas and phases.

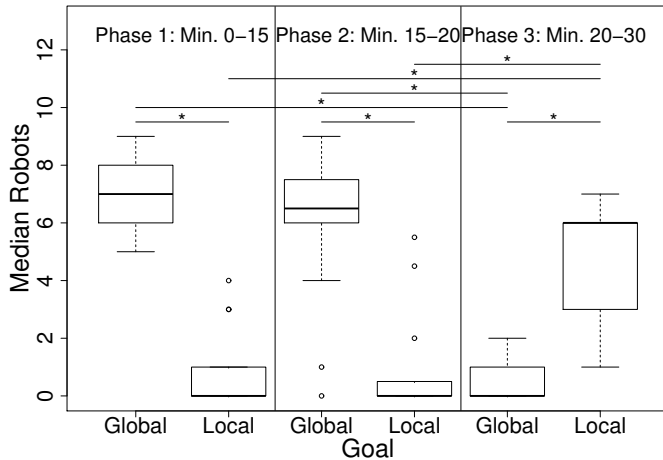


Figure 6: Median number of robots in the respective area over time per phase. Results are significantly different, except in the case of *global optimum* of phase 1 with *global optimum* of phase 2, and also *local optimum* of phase 1 with *local optimum* of phase 2.

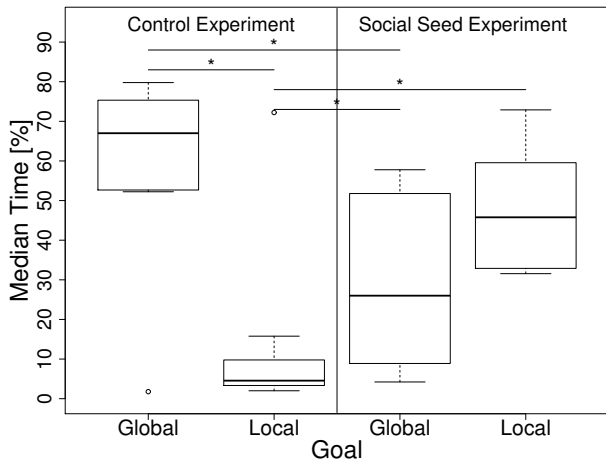


Figure 7: Median time the robots spent in the respective optimum area. Left: Control experiment with 10 robots. Right: Experiment with an immobilized agent placed in the *local optimum* area. Asterisks indicate significant differences with $p < 0.05$.

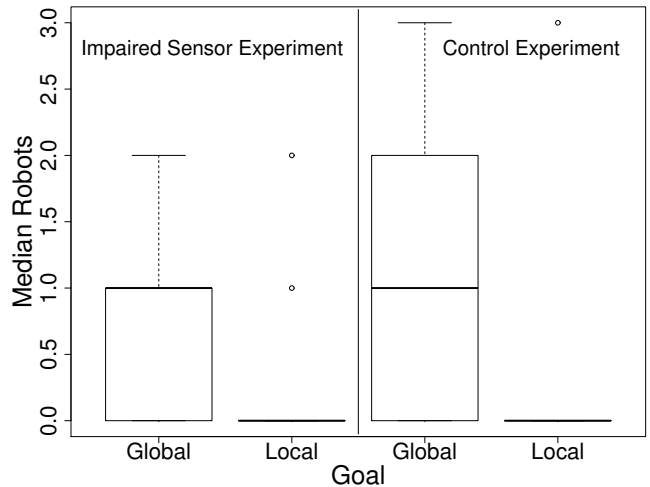


Figure 8: Median number of robots in the two areas. Left: Impaired sensor experiment with 3 fully functional and 7 impaired robots. Right: Control experiment with 3 fully functional robots. There are no significant differences. Significances are tested with the Wilcoxon-Test (*global optimum* against *local optimum*) and the U-Test (optimum of one experiment against the same optimum of another experiment) with $p < 0.05$.

In figure 8 the results with impaired sensor experiments are shown for both the control experiment and the impaired sensor experiment. For the control experiment, the median number of robots in the *global optimum area* is 1 and in the *local optimum area* is 0 (here the swarm consists of 3 fully functional robots). The fact that statistics shows no significant difference between the *global goal area* and the *local goal area* shows, that the aggregation with 3 robots is not very stable (as it was intended in order to keep the chance of increase or decrease of the performance by adding the impaired robots). In the experiment with impaired sensor experiment where the swarm consists of 3 fully functional and 7 impaired robots that cannot measure the temperature (thus having a waiting time of $t = 0s$), the median number of robots in the *global optimum area* is 1 and in the *local optimum area* is 0. Note that in both experiments (control and impaired) only the fully functional robots are counted. The results of the two experiments are not significantly different (tested with U-Test and $p < 0.05$).

4. DISCUSSION

Results of experiment with a dynamic environment in a temperature gradient showed that robots controlled by the BEECLUST algorithm are able to react on changes in the environment reliably. In [17] this experiment was performed in a light gradient and was shown that the robots were able to choose the brightest source of light out of several light spots and that they also were able to react on environmental changes. We performed this experiment in a physically more complex environment. Compared to light, temperature differs in its physical characteristics for example in: warming up period, cooling down, thermal diffusion, turbulences in the airflow. As the experiment lasted 30 minutes, there was

also an air-conditioning necessary so that the temperature of the room (28°C) was stable. Another challenge of working with temperature is the time delay of the measurement. Because of this, after the measurement of temperature was triggered, we had to measure it again a few seconds later and correct the first measurement. Although these physical conditions make experiments more complex, it is shown that robots controlled by the BEECLUST algorithm are still able to react on environmental changes. After switching off the heat lamp of the *global optimum area* the aggregation started to dissolve in the *global optimum area* 2 minutes later and after another 3 minutes robots start to form a cluster in the *local optimum area* (which is now the new global optimum).

The experiment with a social seed was originally designed to analyse the social component of the BEECLUST algorithm. The simulation results [14] showed that the decision-making can be influenced by a social stimulus. In the simulated experiment, the minimum of three social seeds were needed to be placed in the *local optimum area* in order to make the main swarm to spend more time in the *local optimum area* than in the *global optimum area*. In real robot experiments, the effect of social seed is even stronger. Here, a single robot as a social seed is enough to influence the decision making process of the swarm significantly. As the BEECLUST algorithm is derived from the behaviour of young honeybees, experimenting with a temperature gradient gets us closer to the situation that young honeybees are faced with. Because of the strong effect of the social stimuli, we suspect that also in honeybees this social stimuli can be very strong. Further experiments with honeybees can be done to investigate this.

In the last scenario, we study if robots that cannot measure temperature (thus, have always a waiting time of $t = 0$ s) decrease the performance of intact robots. Due to the fact that such impaired robots trigger the measurement of temperature of fully functional robots, there is also the possibility that adding impaired robots could lead to an improvement of performance. Our results show that even adding a big number of impaired robots (more than double the amount of fully functional robots) does not harm the swarms decision making process significantly, although an aggregation of only 3 robots is unstable even without any disturbances. The aggregation stays unstable also in presence of the impaired robots. Thus adding impaired robots does not have any significant effect neither increasing nor decreasing the performance of the main swarm controlled by the BEECLUST algorithm.

5. CONCLUSION & FUTURE WORK

In this paper we study the BEECLUST algorithm with robots under various conditions in a temperature gradient. We conclude that robots controlled by the BEECLUST algorithm are still able to react on environmental changes although the physical conditions in a temperature gradient are more difficult than in a light gradient. When placing a social seed in the *local optimum*, the effect is even stronger than in the simulation results presented in [14], although the social component included in the BEECLUST algorithm is minimal in terms of active communication: namely recognition of another robot. Adding robots with impaired sensors to a swarm of fully functional robots does not harm the swarms' decision making process significantly. This result was very surprising, as robots with impaired sensors do not stop and

therefore cannot be part of an aggregation.

In future we will investigate the application of BEECLUST in relation with other behavioral algorithms. We may also perform the same experiments with real honeybees. As the BEECLUST algorithm is derived from the behaviour of young honeybees, we expect similar effects and thus provide feedback for biological swarm research of honeybees.

6. ACKNOWLEDGMENTS

This work is supported by the following grants: FWF (Austrian Science Fund) "REBODIMENT", no. P23943-N13; EU-ICT "ASSISL bf", no. 601074

7. REFERENCES

- [1] A. Acerbi, D. Marocco, and S. Nolfi. Social facilitation on the development of foraging behaviors in a population of autonomous robots. In F. Almeida e Costa, L. Rocha, E. Costa, I. Harvey, and A. Coutinho, editors, *Advances in Artificial Life*, volume 4648 of *Lecture Notes in Computer Science*, pages 625–634. Springer Berlin Heidelberg, 2007.
- [2] F. Arvin, K. Samsudin, A. R. Ramli, and M. Bekravi. Imitation of honeybee aggregation with collective behavior of swarm robots. *International Journal of Computational Intelligence Systems*, 4(4):739–748, 2011.
- [3] F. Arvin, A. E. Turgut, F. Bazyari, K. B. Arikian, N. Bellotto, and S. Yue. Cue-based aggregation with a mobile robot swarm: a novel fuzzy-based method. *Adaptive Behavior*, 22(3):189–206, 2014.
- [4] F. Arvin, A. E. Turgut, and S. Yue. Fuzzy-based aggregation with a mobile robot swarm. In *Swarm Intelligence (ANTS'12)*, volume 7461 of *Lecture Notes in Computer Science*, pages 346–347, Berlin, 2012. Springer.
- [5] J. Deal. The temperature preferendum of certain insects. *Journal of Animal Ecology*, 10:323–356, 1941.
- [6] ePuck. e-puck desktop mobile robot - website, 2009. <http://www.e-puck.org/>.
- [7] P. Grodzicki and M. Caputa. Social versus individual behaviour: a comparative approach to thermal behaviour of the honeybee (*Apis mellifera* l.) and the american cockroach (*periplaneta americana* l.). *Journal of Insect Physiology*, 51(3):315 – 322, 2005.
- [8] H. Hamann. Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence*, 7(2-3):145–172, 2013.
- [9] H. Heran. Untersuchungen über den Temperatursinn der Honigbiene (*Apis mellifica*) unter besonderer Berücksichtigung der Wahrnehmung strahlender Wärme. *Zeitschrift für vergleichende Physiologie*, 34:179–206, 1952.
- [10] J. M. Hereford. Analysis of a new swarm search algorithm based on trophallaxis. In *Proceedings of the IEEE Congress on Evolutionary Computation, Barcelona, Spain, July 2010*, pages C–7424, 2010.
- [11] J. M. Hereford. Analysis of BEECLUST swarm algorithm. In *Proc. of the IEEE Symposium on Swarm Intelligence (SIS 2011)*, pages 192–198. IEEE, 2011.
- [12] D. Kengyel, H. Hamann, P. Zahadat, G. Radspieler, F. Wotawa, and T. Schmickl. Potential of

- heterogeneity in collective behaviors: A case study on heterogeneous swarms. In *Principles and Practice of Multi-Agent Systems*, PRIMA 15, page 1, 2015.
- [13] D. Kengyel, T. Schmickl, H. Hamann, R. Thenius, and K. Crailsheim. Embodiment of honeybee’s thermotaxis in a mobile robot swarm. In *10th European Conference on Artificial Life (ECAL’09)*, volume 5777/5778 of *LNCIS*. Springer-Verlag, 2011.
- [14] D. Kengyel, R. Thenius, K. Crailsheim, and T. Schmickl. 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.
- [15] S. Kernbach, R. Thenius, O. Kornienko, and T. Schmickl. Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic swarm. *Adaptive Behavior*, 17:237–259, 2009.
- [16] T. Ohtani. Spatial distribution and age-specific thermal reaction of worker honeybee. *Humans and Nature*, 1:11–25, 1992.
- [17] T. Schmickl, R. Thenius, C. Möslinger, G. Radspieler, S. Kernbach, and K. Crailsheim. Get in touch: Cooperative decision making based on robot-to-robot collisions. *Autonomous Agents and Multi-Agent Systems*, 18(1):133–155, February 2008.
- [18] M. Szopek, G. Radspieler, T. Schmickl, R. Thenius, and K. Crailsheim. Recording and tracking of locomotion and clustering behavior in young honeybees (*Apis mellifera*). In A. Spink, M. Ballintijn, N. Bogers, F. Grieco, L. Loijens, L. Noldus, G. Smit, and P. Zimmerman, editors, *Proceedings of Measuring Behavior 2008, Maastricht, August 26-29*, volume 6, page 327, 2008.
- [19] M. Szopek, T. Schmickl, R. Thenius, G. Radspieler, and K. Crailsheim. Dynamics of collective decision making of honeybees in complex temperature fields. *PLoS ONE*, 8(10):e76250, 10 2013.