The influence of environmental circumstances on the emergence of Group and Swarm Behavior

Author:
Thomas Planting
s3030938

Supervisors:
Dr. W.F.G. Haselager
Dr. I.G. Sprinkhuizen-Kuyper

August 31, 2015

Master Thesis
Artificial Intelligence
Faculty of Social Sciences

Radboud University

Abstract

In this project, the evolution of swarm behavior is investigated by means of a computational simulation. The potential evolution of three different types of group behavior, grouping, following and communicating through stigmergic cues, is investigated with different parameter settings in a simulated world. The agents in this world are simple agents directed by a neural network, which will be evolved by means of an evolutionary algorithm. The relationship between the parameters, the starting behavior of the agents and the resulting evolved behavior are analysed from both a biological and a computer science perspective. No significant results are found, but interesting observations are made.

Keywords: Swarm Intelligence, Swarm Behavior, Group Behavior, Cooperation, Neural Network, Evolutionary Algorithm, Simulation, Clustering, Stigmergy, Emergence, Evolution, Environmental Features, Foraging Task

Acknowledgements

Firstly, I would like to thank my supervisors for the freedom I was given to investigate my own ideas, and for the valuable advice and continued support they gave me during the long and often difficult process. Pim, for his relentless efforts to make me channel my wild and often vague and far too broad ideas into something concrete, specific and realistic. Ida, for providing valuable help with the technical aspects of this thesis despite being retired and suffering from her health. Furthermore, I would like to thank all the friends, family and fellow students who helped me get some much needed distraction and relaxation in and around work days, and made this solitary project a bit less solitary.

Contents

Chapter 1

Introduction

1.1 Overview

Swarm Intelligence, or Swarm Behavior, is a field that has fascinated biologists for many years. One definition for Swarm Intelligence is that it is the collective behavior of decentralized, self-organized systems, natural or artificial [Zhang et al., 2013]. Besides its interest to biologists, swarm intelligence is an interesting field for researchers in Artificial Intelligence (AI) or Artificial Life (AL). Swarm Intelligence algorithms tend to be both robust and flexible [Şahin, 2005]. In software development, the property robust refers to the ability of software to deal with problems such as small errors and unexpected inputs. Generally, the robustness of software refers to the ability of the software to consistently perform well even if something went wrong. Swarm Intelligence algorithms tend to be robust because the individuals are each dispensable and the system is decentralized. If a single individual does not function because of some input or something it encountered in the "environment" where it acts, the interactions between the other individuals will only be marginally effected. The flexibility of a system refers to the scope of problems and inputs the system can be applied to, and how much the software will need to be altered for it to work on a different problem or with a different input. Swarm Intelligence algorithms tend be flexible because the coordination between the agents differs according to the environment, because all individuals will respond differently to it.

The robustness and flexibility of Swarm Intelligence algorithms make them interesting from a practical point of view, and the simplicity of the individuals makes it interesting from a theoretical point of view. For example, the foraging behavior of ants has inspired a general purpose optimization technique called Ant Colony Optimization (ACO) [Dorigo et al., 2006]. This technique exploits the relatively low computational complexity of navigation by use of traces that are left in the environment. Generally, the term Swarm Behavior is applied in biology and Swarm Intelligence in computing. I will use them interchangeably in this thesis.

One interesting question about swarms is how they evolved. How can evolution favor individuals with a certain behavior that only functions in a group of likewise individuals? Another interesting question is what kind of swarm behavior would naturally emerge from what kind of situation. For example, what influence does the accessibility of food sources or the presence of predators in an environment have on the characteristics of the swarm behavior that evolves?

The aim of this project is to investigate in a controlled environment in which of a number of environmental circumstances groups or swarms will be able to evolve. Three different types of group behavior will be investigated. This research aims to uncover parameters that have an effect on the emergence of these kinds of group behavior, and what kind of influence these parameters have on this behavior. It will be done by means of a simulation. An evolutionary algorithm will be employed. This research will deal with individuals, or agents, which only follow a few simple behavioral rules. Additional aims of this thesis are to investigate whether these different kinds of group behavior emerge from each other, and what effect a select few parameters have on the emergence of these behaviors. This is exploratory research, statistical tests will be performed, but the data will also be analysed in a less rigid way, and any observed trends will be mentioned and discussed. The behaviors that will be investigated are the following:

1. Clustering: whether agents tend to cluster in groups. If significantly more or several groups of agents can be distinguished by a clustering algorithm at the end of a simulation run, the agents will be considered to be clustering.

2. Following: whether agents tend to actively move to other agents. If agents move towards other agents significantly more than the baseline, the agents are considered to be displaying following behavior.

3. Stigmergy: Stigmergy is the phenomenon of indirect communication mediated by modifications of the environment Marsh and Onof [2008]. In this project, the agents will leave behind trails. This is hard-coded in the environment and is not

what is being investigated. What is being investigated is if other agents will respond to these traces. If agents are influenced by the behavior of other agents through the stigmergic trails they leave behind, they will be displaying stigmergic behavior. Since the leaving behind and following of stigmergic trails can lead to relatively complex and effective decentralised behavior, this is considered to be swarm behavior. See Section 1.2 for more information on stigmergy and how it can constitute swarm behavior.

Note that of these three behaviors, only stigmergy can be considered to be proper swarm behavior. Clustering and Following are merely group behavior, but could potentially lead to the emergence of swarm behavior. This will be discussed in Section 1.3. The parameters that will be investigated are the following:

- Predators: Whether there are predators present in the environment. This variable can be either on or off. If it is on, the environment will be populated with predators which will hunt the agents.
- Group Protection: If predators are present, the group protection variable determines whether predators will cease moving to agents if there are multiple agents together. If the Group Protection parameter is on and a number of agents are close to each other, they will not be followed by predators.
- Food Clustering: In the default setting, the food will be distributed randomly across the world. With clustering active, groups of food will be spawned in clusters of a pre-determined size with a maximum distance between every 2 food sources.

1.1.1 Research Questions

Research Question 1:

- With which parameter settings will clustering behavior emerge from novel agents?
- Hypothesis: with food clustering, predators and group protection active. I expect predators and group protection to be necessary to provide an incentive to form groups. I expect food clustering to be necessary to prevent groups from depleting their food sources too quickly.

Research Question 2:

- With which parameter settings will following behavior emerge in the simulation from novel agents?
- Hypothesis: with food clustering active. I expect the agents will follow other agents to find clusters of food.

Research Question 3:

- With which parameter settings will stigmergic behavior emerge in the simulation from novel agents?
- Hypothesis: with food clustering active. I expect the novel agents will first learn to follow each other and later learn to follow stigmergic trails. Once agents have already learned to follow other agents, the following of stigmergic cues can provide a more reliable mechanism to achieve the same goal, the finding of food clusters.

The fourth and fifth Research Questions are optional and can only be answered if a significant effect is found for the first or second Research Questions respectively, since they require agents which have already developed a behavior.

Research Question 4:

- With which parameter settings will stigmergic behavior emerge in the simulation from following agents?
- Hypothesis: with food clustering active.

Research Question 5:

- With which parameter settings will will stigmergic behavior emerge in the simulation from clustering agent?
- Hypothesis: with food clustering active. If the clustering agents are not already displaying following behavior, I expect the agents to first learn to follow.

The following Research Question is optional and can only be answered if a significant effect is found for any of the three operationalisation values for any of the settings. This is a requirement for this test since it is meaningless to test for the effect of parameters on the investigated behaviors if said behaviors can not be shown to emerge in any of the settings.

Research Question 6:

- What effect do the different parameter settings have on the emergence of the different behaviors?
- Hypothesis: Food Clustering will have a positive effect on all three behaviors, and predators and group protection will have a positive effect on clustering.

1.1.2 Scientific Relevance

The scientific relevance of this thesis is twofold. Firstly, it aims to provide insight into the nature of the evolution of swarms, and the circumstances under which certain kinds of swarms can evolve. This will be interesting for evolutionary biologists. There has already been extensive research into each of the three behaviors that are investigated in this thesis (clustering, following and stigmergy), but how these behaviors may emerge from each other has not seen much research. Furthermore, most biological research is focused on specific species and empirical in nature. This research focuses on finding general trends, making use of abstract and virtually simulated environments and agents. The general and abstract point of view adopted in this thesis, though not unique, certainly distinguishes it from most existing research. This thesis could provide a direction for future research into the evolutionary steps of group and specifically swarm behavior.

Secondly, although this thesis will only cover a limited amount of behavior rules and a specific environment and task, it may provide valuable insight into the relationship between behavior rules of swarm individuals and the emergent behavior of a swarm. This can be valuable for AI researchers, as emergent behavior is often difficult to deduce from the simple rules of the individual elements, and has not seen extensive research yet. Most existing AI research into Swarm Intelligence exploits a specific, known Swarm Intelligence algorithm in a certain niche.

1.2 Background

Conform the definition by Zhang et al. [2013], Swarm Intelligence is the collective behavior of decentralized, self-organized systems, natural or artificial. Instead of conscious deliberate coordination between individuals, all individuals perform their own actions and the collective behavior naturally *emerges* from their interactions. The concept of emergence is a crucial one and deserves additional explanation. One way to describe emergence is through the following quote: "emergence relates to phenomena that arise from and depend on some more basic phenomena yet are simultaneously autonomous from that base" [Bedau and Humphreys, 2008]. More basic phenomena interact in a way that causes the more complicated phenomena to emerge. Throughout this thesis the terms "simple" or "basic" and "complicated" or "higher-order" refer to the complexity of the rules that describe a certain phenomenon/behavior with respect to the environmental situation of the object or individual. An example of such a simple behavioral rule is if you smell food, move in the direction the smell comes from. These terms will generally only be used with relation to each other, in which case the "complex" phenomenon has a higher order of complexity than the "simple" or "basic" phenomenon. In the quote above, a complex phenomenon arises from basic phenomena. Another key point that was highlighted in the quote above is that the arising phenomena are autonomous from the basic phenomena that cause them. This may appear to be a contradiction, but it means that the emergent phenomenon is not merely a sum of the basic phenomena but something above them which arises from it. An example of emergence is cognition. All neurons of the brain exhibit simple input-output mappings, but the general scientific view is that together they are responsible for human cognition, including dreams, thoughts, feelings etc. This human cognition is very different from the behavior of individual neurons. When we apply the concept of emergence on living entities, we often refer to it as "emergent behavior". Another way to describe emergence or emergent behavior is through the following quote: "A reasonable way of thinking about emergent behavior might be to focus on the level or scale at which the rules reside. If the rules are specified at a low level, for example, the individual termites, and the patterns and structures, like termite mounds, emerge at a scale where there are no rules specified, we may call this emergent behavior" [National Research Council, 2008]. Note that this is consistent with the previous quote, emergence arises from more basic phenomena, thus the rules of the system are specified at a low level. From these rules, the higher level phenomenon, which is autonomous to them, emerges.

Swarm behavior is emergent behavior, the behavior of the entire swarm emerges from the behavior of the individual members of the swarm. This means swarms are highly decentralized, there is no central coordination in a swarm. Instead, swarm behavior is emergent behavior arising from simple behavioral rules that do not involve any central communication. An example of such a simple behavioral rule within the context of swarms is the following: if you smell food, move in the direction the smell comes from. Swarm behavior emerges from locally operating agents, none of the agents have an internal representation of the entire environment, they only interact locally on their direct inputs.

Often, the collective behavior of swarms can solve problems which the individuals would be unable to solve on their own. A defining characteristic of swarms is that the behavior that emerges from their interactions is able to solve relatively complex problems, which the behavior of the individuals cannot solve. This is conform the idea that emergence is the arising of complicated phenomena from simple phenomena. An example of this is an ant colony. Collectively, ant colonies are able to solve complex problems. However, as said by Garnier et al. [2007], "surprisingly, the complexity of these collective behaviors and structures does not reflect at all the relative simplicity of the individual behaviors of an insect". An example of such a complicated problem that ant colonies are able to solve but individual ants are not is pathfinding. As a colony, ants are able to locate the shortest path from the nest to a food source [Goss et al., 1989]. For some ant species, when ants move, they leave behind a chemical substance called pheromones, which attracts other ants. If multiple ants move to a food source, the ants that took the shortest route are more likely to have returned to the nest first. These ants will likely take the same path back to the nest, because they marked it with pheromones. When these ants return to the nest, the shortest path is thus slightly more marked with pheromones, and thus more attractive to ants that leave the nest to gather food. This will cause other ants to follow the path and strengthen the pheromone path again. Moreover, the ants that take the shorter route will reinforce their pheromone trail faster than the ants that take the longer route, because the shorter route is passed more often in the same time. By these kinds of positive feedback loops, swarms are able to solve complicated problems despite the individuals all being relatively simple. As stated by Garnier et al. [2007], "The colony "as a whole" is able to produce an efficient collective response that far exceeds the scale and abilities of a single individual ant." The collective behavior of the swarms is able to solve problems that are too complex for the relatively simple ants too solve individually. Moreover, the individual ants do not attempt to solve the problem individually but each perform their own basic tasks, from which the behavior that solves the complex task emerges. This example is highly relevant in this research, because the following of stigmergic trails is one of the behaviors that is being investigated, and the only investigated behavior that is considered to be swarm behavior.

An important feature for many swarms, especially for lower cognition "insect like" swarms that will be the main focus of this thesis, is stigmergy. In its most general form, stigmergy is "the phenomenon of indirect communication mediated by modifications of the environment" [Marsh and Onof, 2008]. It is a process via which unorganized behavior of individuals serves as stimuli for the actions of other individuals by leaving traces in the environment. This way, complex behavior can emerge from seemingly unorganized individuals. Stigmergy enables members of swarms to act on information that is not directly perceivable by them, by using the cues about the environment left by others as input. The above example about the finding of the shortest path by ant colonies is a classic example of stigmergy. In that example, the thickest pheromone trail can be viewed as representing the shortest path from the nest to the food, which the ants use to quickly bring the food to the nest. An ant who just leaves the nest and can not perceive the food will still be able to follow the pheromone trail to it, thus acting on information that is not directly perceivable by it.

Another important characteristic of many swarms is that the individuals only function if other individuals follow certain behavioral rules as well. For example, if only one ant would drop and follow pheromones, the shortest distance to a food source would not be found. The pheromones would only cause the ant to keep following the same path, even if it is an inefficient path, because there is no competition between different pheromone trails. It would appear that some swarm behavior only functions if individuals can count on others to also exhibit the right behavior.

1.3 The Emergence of Swarms

The question, then, is how can this behavior have evolved initially? How can a certain behavior evolve which does not function if group members do not exhibit that behavior as well? Answering this question is one of the main aims of this thesis, and will be discussed more in depth later on. Swarm Behavior can be considered to be a special case of group behavior. In this Section, I will first discuss some theories about the emergence of group behavior in general. After that, I will narrow it down to swarm behavior and discuss which of these theories can be applicable for that and what else might be needed for true swarm behavior to emerge.

1.3.1 Group Behavior

There has already been extensive research into the emergence of group behavior in general. The emergence of group behavior can be split up between cooperation, in which the individuals actively aid each other, and profiting, in which the individuals merely profit from the actions of others. In this Section, an overview of some of the most prominent theories will be provided. After that, I will discuss how Swarm Intelligence differs from more "general" group behavior, and what this means for the emergence of Swarm Intelligence.

Cooperation

Cooperation is commonly defined as any adaptation that has evolved, at least in part, to increase the reproductive success of the actor's social partners [Ross-Gillespie et al., 2007]. The problematic feature with cooperation is that to it may be costly for an individual to help another, with no direct benefit to the individual itself. This phenomenon is called altruism [Trivers, 1971]. In evolutionary biology an organism is generally said to be altruistic when its behavior benefits other organisms but has a cost to itself [Kerr et al., 2004]. There are several theories for the evolution of cooperation and altruism, which are not mutually exclusive and may work in unison. Some of the most prominent of these theories are:

- Kin Selection: If individuals aid other individuals which are related to them, the genes that cause altruism will pass on [Hamilton, 1964].
- Direct Reciprocity: If there are repeated encounters between the same individuals, they may increase the chance for the other one to cooperate by cooperating themselves [Nowak, 2006].
- Indirect Reciprocity: By cooperating, an individual might build a reputation for itself, which can be rewarded by others. By not cooperating, other individuals may punish it [Boyd and Richerson, 1989].
- Spatial Reciprocity: If the success of an individual is partially dependent on the success of the individuals in the same group or geographical area, it will be in the individual's benefit to aid them [Nowak and May, 1992].
- Group Selection: If groups of individuals who help each other perform better than others, those groups will have an evolutionary advantage over groups that do not [Traulsen and Nowak, 2006].

Profiting

Profiting is the phenomenon of individuals forming groups to increase their own reproductive success. Individuals that somehow manage to profit from the behavior of others perform better. Only individuals who profit survive, and in the end they all "profit" from each other, resulting in cooperation. This is much more likely if the sum of the fitnesses of the group increases if individuals "profit" from each other. Profiting stands in contrast with "true" cooperation, where individuals perform altruistic acts to their group members.

An example of profiting is the following: if there are large food sources in a world (for example dead animals for small scavengers), it will be beneficial for single scavengers to follow others if they can sense whether they have picked up a trail. Such scavengers will then have a higher chance to survive, and thus have more chance to reproduce. This will give an evolutionary edge to their genes, and future generations of these scavengers may all follow each other if they head towards food. Another example of profiting is the forming of large groups of individuals to discourage predators from attacking them. Profiting does not exclude the future possibility of cooperation within groups. Groups of individuals who initially band in groups to profit from each other's behavior might later advance into "true" cooperative behavior through any of the mechanisms described above.

1.3.2 Swarms

The difference between regular group behavior and swarm behavior is that something else is needed for swarm behavior. Conform the definition by Zhang et al. [2013] , for a certain group behavior to be swarm behavior, there has to be a coherent collective behavior that naturally emerges from the interactions of the individuals.

Broadly speaking, there are two different ways in which swarms could possibly emerge: from individuals and from groups. If they emerge from individuals, it means these individuals will form groups in a way in which a collective behavior results from their interactions. If they emerge from groups, it means the individuals will first form groups through the means described in the previous subsection. After that, the resulting groups will evolve a collective behavior that emerges from their interactions. I will first discuss the possible ways in which the kind of individuals that could later evolve into swarm individuals can form groups. After that, I will theorize about how individuals or groups could possibly form swarms

From Individuals to Groups

In general, groups that will later evolve into swarms do not necessarily have to be different from other groups. Therefore, the mechanisms that cause those groups to emerge do not have to diverge from mechanisms described above. Since many swarms consist of simple individuals, however, it is safe to rule out processes which require higher cognitive functioning as a basis for group forming. It is still possible that said processes play a role in the formation of swarms of animals with a higher cognition, but those will not be the focus of this thesis. Another reason not to focus on mechanisms that require higher order cognitive abilities is that swarm mechanics tend to be very simple. Thus, if the same mechanics that play a role in the swarm behavior played a role in the formation into groups of the ancestors of those swarm individuals, they are most likely not caused by mechanisms which require higher order processing.

An important feature that is present in many swarms is Stigmergy. If the groups that will be formed will later evolve into stigmatic swarms, which communicate through stigmatic cues, it is likely stigmergy already plays a role in the group behavior. If this is the case, the emergence of stigmergy in groups will have to be explained.The plausibility of the different theories with respect to swarms will be discussed briefly below:

- Kin Selection: Possible, kin selection has nothing to do with the individual's cognitive complexity.
- Direct Reciprocity: Implausible, depends on the individuals to remember who helped them.

- Indirect Reciprocity: Implausible, depends on the individuals to remember other individual's reputation.
- Spatial Reciprocity: Possible, spatial reciprocity has nothing to do with the individual's cognitive complexity.
- Group Selection: Possible, group selection has nothing to do with the individual's cognitive complexity.

From Individuals and groups to Swarms

An other way to explain the emergence of swarms is through stigmergy and the profiting theory. If individually operating (non-swarm) individuals perform actions in certain steps, which change the environment, they could use their own previous actions as clues for what to do next. If another individual meets the environment in the state it reaches after the first few "steps" of work, it does not matter that it did not do the first steps of the work itself, and it may just continue as if it did. This way, multiple individuals would start cooperating. This behavior could easily provide an advantage because it allows for sequential operations without requiring a sense of memory in the individuals. If this cooperation increases the sum of the utilities of the involved individuals, evolution would favor them. This explanation comes forth from profiting, but from there the stigmergic group could evolve into higher order cooperation through any of the other theories.

With this example, however, the swarm agents would still function as individuals without their groups. For many swarms, this is not the case. The question that remains is how behavior that does not function on its own could have evolved initially. Except for genetic drift, which is not a satisfying explanation given the abundance of swarms in our world, behavior can only evolve if it is already profitable for the individuals themselves, their kin or their groups. It appears probable that the swarm behavior is reached by means of many intermediate steps. After all individuals have passed such an intermediate step, they can count on others to exhibit a certain behavior and another intermediate step can be taken next. This way, complicated behavior that is dependent on the behavior of others could evolve step by step. These steps can initially be caused by random mutations, or serve a different purpose. In the latter case, exaptations would play a role. With the pheromone example, this would mean that individuals would initially start dropping pheromones as a result of a random mutation or to achieve some other goal, to dispose of waste products for example. When other individuals evolve to start following these pheromones and thereby increase their fitnesses, however, dropping pheromones will be beneficial for the group and an individual's own genes, and thus be favored by evolution. That way, more sophisticated cooperation can emerge over time. In reality, this is most likely much more complicated, and there are likely to be other factors in effect at the same time. It does, however, provide a basic explanation for the emergence of swarm behavior. It is likely that this kind of cooperation will be easier to achieve from individuals who are performing basic group behavior than for solitary individuals. Individuals who are already cooperating in a basic way will generally be closer to each other, and will be more likely to pick up cues for swarm behavior. For example, if individuals have already learned to follow each other, they will have a lot of interaction with the stigmergic trails of other agents, making it more likely to develop a following behavior for the trails.

Chapter 2

Methods

2.1 Research Plan

2.1.1 The Aim

This research aims to uncover parameters that have an effect on the emergence of several kinds of group or swarm behavior, and what kind of influence these parameters have on this behavior. The research will be done by means of a virtual simulation, which will be explained in detail in Section 2.3. The research focuses on the emergence of clustering, following and stigmergic following behavior, which will be tested by means of separate simulation runs. Each of these behaviors will be defined in the context of the simulation in Section 2.1.7. The evolution of stigmergy is investigated both from novel agents and with agents which have already learned following or clustering steps, provided those agents exist. If any of the behaviors is found to be present after the evolution, the influence of the parameters on the emergence of these behaviors is analyzed.

2.1.2 The Simulations

This section will detail the simulations that will be run and their corresponding parameter setups. A total of 4 experiments will potentially be run, which can be grouped into two categories: experiments that work from novel agents and experiments that work from agents that have already learned one of the investigated behaviors. In the first category, one experiment will be run with novel agents to investigate the potential emergence of clustering and following behavior, and one will be run to investigate the potential emergence of stigmergic behavior. These two experiments are separate to be able to analyze the differences between group forming with and without the presence of stigmergic traces. The first experiment will investigate Research Questions 1 and 2, and the third experiment will investigate Research Question 3 (see Section 1.1). In the second category of experiments, the emergence of Stigmergic behavior will be investigated from agents that have already learned to cluster and from agents that have already learned to follow. These experiments will only be run if any significant results are found in experiment 1. These experiments will investigate research questions 4 and 5.

For all three behaviors, various parameters have been identified which could be relevant for the emergence of the behavior. Separate simulations will be run for each setting of the parameters that are investigated. For each parameter setting, 10 train simulations will be run to provide 10 different groups of genotypes. For each of these genotypes, 10 single-generation test simulations will be run, resulting in 100 test runs per setting. Each test run will be given a single value which represents how strongly the researched behavior is considered to be present. The measures for this will be discussed in the Section Operationalization. These test simulations will be compared with an equal number of baseline test simulations. The baseline test simulations are run from the same gene pool, but during the baseline tests the agents are unable to perceive stigmergic traces or other agents. These restriction in the sensor inputs of the agents are in place to ensure the agents do not perform group behavior in the baseline tests. The baseline tests will also be given an operationalization value. The agents in a setting will be considered to be displaying a certain kind of group behavior (clustering, following or stigmergy) if the operationalisation values of that behavior are significantly higher in the experimental condition than in the baseline condition. After this, if any significant results are found for the operationalisation values, a sensitivity analysis will be performed on all simulations where the investigated behavior was found to be present. After this, new test simulations will be run with the genotypes of each setup where one of the group behaviors was found to be present in a scarce world to compare their performances. All simulations will be run with the same Random Seed, which makes them exactly replicable.

2.1.3 Experiment 1: Clustering and Following

The values of the parameters that will be varied for the clustering and following behaviors are shown in Table

2.1.

Table 2.1: Parameter Values

Parameter	Values
Predators	Off/On
Group Protection	Off/10/5
Food Clustering	Off/10/50

10 simulations will be run for each level of food clustering with the predator parameter Off, and with the predators parameter On. If the predators parameter is off, the Group Protection parameter will be irrelevant and be set to a default of off. If the predator parameter is on, all 3 levels of Group Protection will also be varied. This makes for a total of $(3+3*3)*10 = 120$ simulations for this behavior. The resulting data will be analyzed on both clustering and following behavior.

2.1.4 Experiments 2-4

For these experiments, only the food clustering behavior will be investigated, which will take on the same 3 values as in Experiment 1, as shown in Table 2.1. The other two parameters will not be investigated for these experiments to save computation and disc space, and because they are not expected to have any influence on the potential emergence of stigmergic behavior. The Predators and Group Protection parameters are expected to have an influence on the emergence of Clustering behavior, but not on stigmergic behavior. These simulations are ran with novel agents (Experiment 2), agents which have already learned to cluster (Experiment 3) and agents which have already learned to follow (Experiment 4). Experiment 3 will only be run if the agents have learned to cluster in at least one setting in experiment 1, and experiment 4 will only be run if the agents have learned to follow in at least one setting in experiment 1. Each of these three experiments will run 10 simulations for all 3 values for the Food Clustering parameter, which makes for $3 * 10 = 30$ simulations per experiment. The resulting data will be analysed on stigmergic following behavior.

2.1.5 Comparison

After the separate tests for the different behaviors, if any of the researched behaviors is found to be present, a new sequence of test simulations will be run with the genotypes that resulted from the successful initial simulations. In these simulations, the different genotypes will be compared to each other in terms of average utility on the same worlds. The worlds where the agents will be compared to each other are the same worlds from the settings where the successful genotypes come from, as well as the basic setting. The exception to this is the amount of units of food, this will only be half the original amount. In each of these settings, 100 baseline and experimental test simulations will be run with each of the successful genotypes. The fitnesses of the agents at the end of the test simulations will be compared with each other. The purpose of this comparison is to determine whether the existence of the behaviors that are being investigated makes the agents more successful in a scarce world than agents who are not displaying those behaviors, and to compare the different layers of group behavior to each other (grouping, following and stigmergy).

2.1.6 Sensitivity Analysis

For any settings with which the investigated behavior is found to be present, a sensitivity analysis will be conducted. New simulations will be run for each such setting with a number of basic parameters,which will henceforth also be referred to as the validation parameters, tweaked. The values for these parameters are found in table 2.2.

Table 2.2: tab:Basic Parameter Variations

Number of Agents	10/25/50
Number of Food Sources	20/50/100
Field Size	50/100/200
Number of Nests	1/2/3
Obstacles	0/5/20

For all settings which are being validated, 10 simulations are run for each combination of the validation parameters. The resulting genotypes are tested through 10 single-generation runs in the same way as the initial simulations. If any settings are not validated, the specific validation parameter values for which it worked and those for which it did not work will be reported and analyzed.

2.1.7 Operationalization

This section details quantitative methods to identify each of the behaviors that are being researched.

Clustering

To determine whether the agents can be considered to be clustering, the positional information of the agents at the end of each test simulation is saved. The agents will be represented as 2 dimensional data points, each will have their x and y positions as values. From this information, groups will algorithmically be identified and subsequently evaluated. This method operates under the assumption that if there are clear groups present, the algorithm will yield the correct groups. Considering the low dimensionality of the data, this appears to be a reasonable assumption.

To identify groups, a clustering algorithm will be applied. In Data Science, the concept of clustering is well known and well researched. It can be defined as "the unsupervised classification of patterns (observations, data items, or feature vectors) into groups (clusters)" [Jain et al., 1999]. To cluster the agents, the Density Based clustering technique DBSCAN will be applied. This is an older clustering technique which still enjoys widespread use [Tan et al., 2006, Chapter 5]. This clustering technique was chosen because it is able to handle data with an unknown number of clusters, quoting Shah et al. [2012]: "One of the advantages of using these techniques is that [the] method does not require the number of clusters to be given a prior". DBSCAN is also able to find arbitrarily shaped clusters. Both of these features are important for finding clusters of agents in this research. The basic functioning of the DBSCAN algorithm is as follows:

- Each point p that has at least k points within a range of d is considered a core point.
- A point q is said to be directly reachable from p if p is a core point and the distance between p and q is smaller or equal to d.
- A point q is considered to be reachable from p if there is a path $p_1...p_n$ with $p_1 = p$ and $p_n = q$ where each p_{i+1} is directly reachable from p_i.
- Each group of core points that are reachable from each other and each non-core point that is reachable from these core points is considered to be a single cluster.
- Each non-core point that is not reachable from another point is not considered to be part of a cluster.

In the operationalization, the following parameters were chosen:

- $k = 4$ this 4 is considered a "reasonable" k for 2-dimensional data by [Tan et al., 2006, Chapter 5].
- $d = proximityRange = 10$ Agents that can detect each other can deliberately move to each other, and are considered to be in the same group
- The distance measure is the Euclidean distance with respect to the x and y positions of the agents. The distance between two agents a_1 and a_2 is defined as $\sqrt{(a_{1x} - a_{2x})^2 + (a_{1y} - a_{2y})^2}$.

For more information on the DBSCAN algorithm, see [Tan et al., 2006, Chapter 5]. When potential clusters have been identified, a value will need to be assigned to the simulation, which represents in how far the agents can be considered to be clustering. Various methods exist to evaluate clusters in Data Science, but most of these have as goal to determine if the current clustering of the data is the right one, if it is better than other clusterings of the data. The clustering algorithm will always come up with a solution, but this does not mean the agents can actually be considered to be clustering. The question in this research is *whether natural groups exist in the data at all*. There are still several methods for determining this in Data Science, but the purpose of those methods is different from the purpose of this operationalization. Existing methods measure *how clearly the data is clustering*. If a dataset can clearly be separated in several different groups, a high score will be given to that dataset. Typically, measures like within-group distance, the average distance between data points which belong to the same group, and between-group distance, the average distance between data points that belong to different groups, are used to evaluate how clearly natural clusters are present in a data set. In this research, the purpose of the cluster evaluation is different: to determine *how much the data is clustering*. How easily the different clusters of agents can be distinguished (determined by between-group distance) from each other is not interesting for this research, nor is how close agents within a group are to each other (determined by within-group distance). If several groups are close to each other, this does not mean the agents should be considered to perform less like a group than if those groups were far apart from each other. Likewise, agents who are relatively far away but still within proximity range of each other do not count less as a group than agents who are close to each other. To determine how much the agents are clustering, we want to know how many agents are in clusters, and how big these clusters are. The measure that is used in this research is the following formula: $o = 1/(c + n)$, where o is the operationalisation value for clustering, c is the total number of clusters found by DBSCAN and n is the number of agents that were not assigned to a cluster. This gives a higher score for more agents being in clusters and for fewer but bigger clusters.

Following

To determine whether the agents are displaying following behavior, the actions of the agents are measured. Each time an agent moves towards any other agent with the action "move to agent" or "move to high fitness agent", 1 is added to a global variable f. This is done for every agent in every test simulation run. The actions the agents can take will be explained in Section 2.3.5. The f value at the end of a test simulation is the operationalisation value for the following behavior.

Stigmergy

To determine if Stigmergic trails are being followed by the agents, the difference in the concentration of stigmergic traces in the environment at the end of each test simulation will be measured. Since every agent leaves a stigmergic trail behind, the more agents follow a trail the stronger the stigmergic concentration becomes. Therefore, a large difference of stigmergic concentrations on different parts of the environment is an indication that stigmergic cues are being followed. The standard deviation of the values of these squares is taken as the measurement for the differences in concentration. This value will henceforth be referred to as s. A high standard deviation means there are large differences in concentration.

2.1.8 Statistical Tests

The agents in a setting will be considered to be displaying a certain kind of group behavior (clustering, following or stigmergy) if the operationalisation values of that behavior are significantly higher in the experimental condition than in the baseline condition. Since the tests are conceptually distinct, and only deal with with a single dependent variable (the operationalisation value of the group behavior that is being investigated with that test) and two conditions (baseline and experimental), the data is analysed with a series of t-tests. A one-tailed t-test will be performed for each setting for each group behavior that is being tested for in that setting. The alternative hypotheses for these tests are that the values that are being compared is higher in the experimental condition than in the baseline condition. In these t-tests, the operationalisation values of each test run are considered single observations. Since 10 train simulations will be run for each setting, and 10 baseline and experimental test runs will be run for each train simulation, this makes for a total of 100 data points for both conditions in the t-tests. On top of that, a one-tailed t-test will be run for each setting to test whether the fitness average in the experimental condition is significantly greater than in the baseline condition. For these tests, the average fitness of a test simulation is treated as a single observation in the same way as the operationalisation value in the t-tests that were mentioned above. The purpose of this test is to investigate for settings where group behavior has emerged whether this group behavior has improved the performance of the agents.

2.2 The Parameters

In this Section, I will discuss the Parameters that will be varied during the simulations. These parameters can be divided into two types: the experimental Parameters and the validation Parameters. The experimental Parameters are the object of this study and their values define the setting a simulation is set in. The validation parameters are parameters that are not directly relevant for the research but will be varied between the simulations to test the robustness of the experimental setup.

The experimental parameters

During the simulation, several factors are varied to determine the possible effects they have on whether swarm behavior will emerge in the simulations, or the effect on the kind of swarm behavior that will emerge. The factors that are varied are the following:

Predators

Whether predators are present in the environment. This variable can be either on or off. If it is on, the environment is populated with predators which behave as described in Section 2.3.3. The default number of predators and the predator's vision range are described in the standard settings Table 2.3.

Group Protection

If predators are present, the group protection variable determines whether predators will cease moving to agents if a certain number of agents are within sensory range of each other. This variable can be either off or on, and if it is on it can take on several values, which represent the number of agents that need to be together to cause the predators to stop moving to the agents.

Food Clustering

Another thing that is varied is the clustering of food. In the default setting, the food is distributed randomly across the world, which can be seen as a having food clusters of size 1. With clustering active, groups of food are spawned in clusters with a minimum food density of 1 unit of food per 5 squares.

Energy decay

The amount of energy the agents automatically lose each action, besides the energy they lose for performing actions.

Stigmergic trace possibilities

This parameter determines whether the agents are able to detect stigmergic traces in the environment. It can be either on or off.

2.2.1 The validation parameters

Besides the experimental parameters, many other parameters can be adjusted. To better ensure the generalisability of this research, a sensitivity analysis will be performed if any significant effects are fond. "A "sensitivity analysis" of these parameters is not only critical to model validation but also serves to guide future research efforts."[Hamby, 1994] A number of parameters have been identified and will be adjusted slightly during several experimental runs for settings for which significant effects were found.

Number of Agents

The number of agents in the environment.

The Abundance of Food

The number of food sources in the environment.

The Field Size

The size of the simulation world are varied as a parameter. The height is always equal to the width, and the number of squares in each direction is considered the size e.g a field of size 100 will be a 100 by 100 field. Varying this parameter changes the density of the objects in the field, which could be relevant for following and foraging behavior.

The number of nests

The default simulations are run with only a single nest. In the sensitivity analysis the effects of having several nests are investigated.

2.2.2 Obstacles

In some validation simulation, impassable obstacles are present in the environment.

2.3 The Model

A model was built in the Java programming language. The environment consists of a virtual 2D grid world of a fixed size of 100 by 100 squares. The environment is enclosed on all sides by impassable walls. A picture of the model is shown in Figure 2.1. The simulation features a food gathering task. The environment is populated with agents who slowly lose energy over time and by performing actions. To replenish their energy, the agents need to consume food. In some simulations there are also predators present, which the agents need to avoid. The simulation features a set number of generations, over which the agents evolve. During a generation, no agents are removed. At the end of each generation every agent has a chance to be selected as a parent for the next generation based on its fitness. The fitness of an agent is based on its energy level and on whether it has been caught by a predator. The evolutionary mechanics is explained in more detail in the Section The Evolutionary Dynamics. The environment is populated with the following objects: Food Sources, Nests, Predators, Obstacles and Agents.

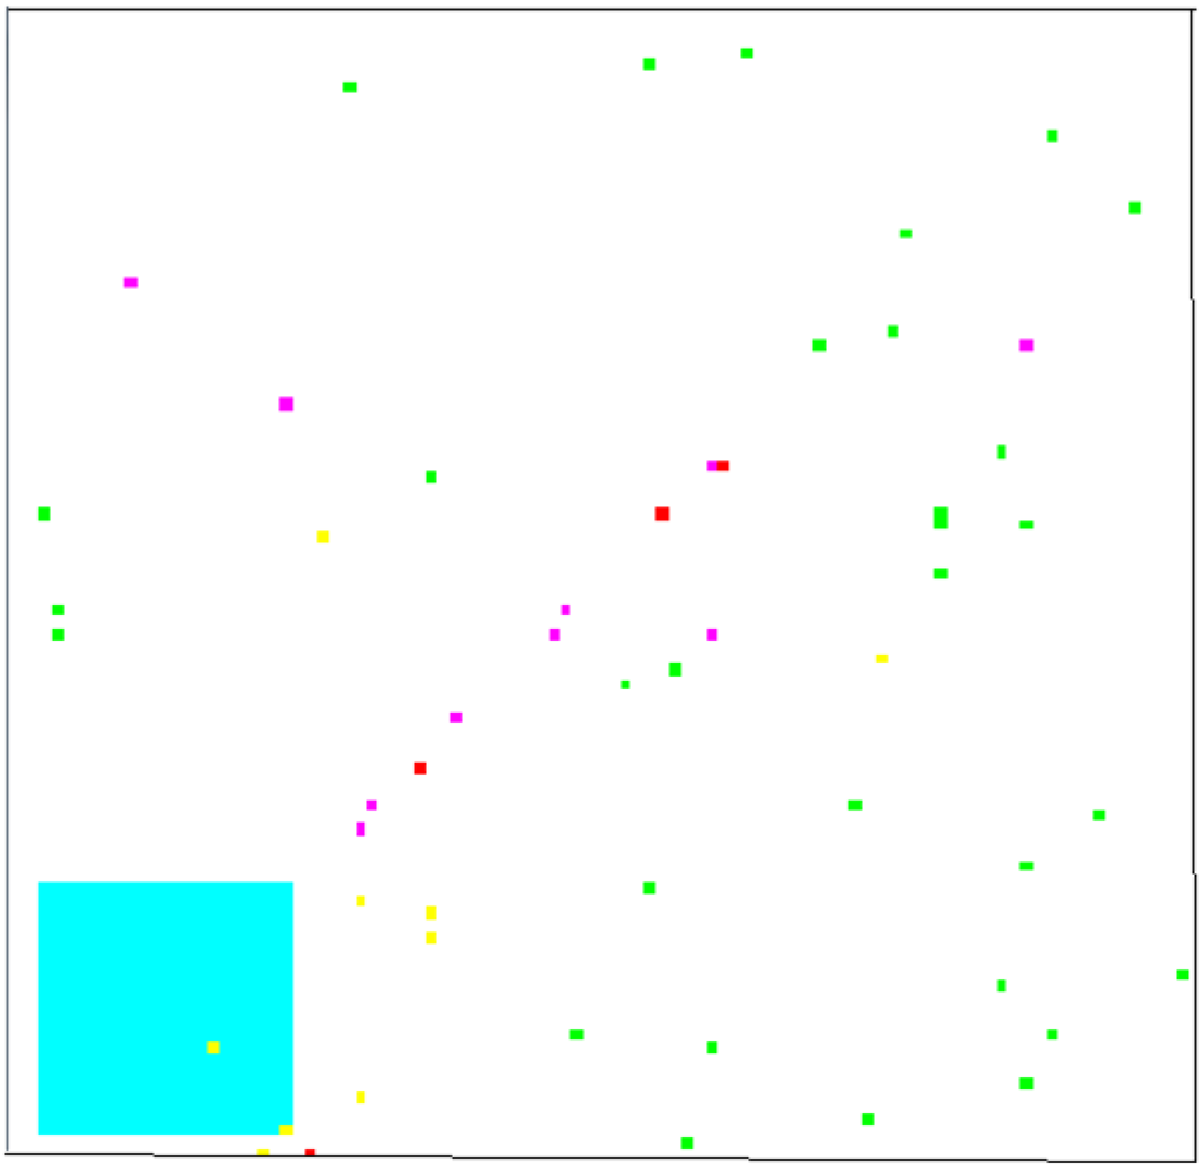

Figure 2.1: A visualisation of the environment during a simulation run. The light blue square is a nest, the green squares are food sources, the yellow squares are agents, the pink squares are agents who are carrying food and the red squares are predators.

2.3.1 Food Sources

To replenish energy, Agents need to consume food. In some simulations the food is scattered throughout the world, and in some worlds it is clustered (depending on the parameter settings). When a food source is depleted it is removed from the world. The food can be located by agents who are within a certain range of it. At any time, a maximum of 50 units of food can be present in the environment, and half of that is present at the start of the simulation. The food replenishes over the course of a simulation. The amount of food that replenishes depends on how much food there is present compared to the maximum amount food. At each time step, the amount of food that is replenished equals half the difference between the current amount of food that is present in the environment and the maximum amount of food. This replenished food is added at random positions (or, in simulation runs where food clustering is active, in random clusters) in the environment.

2.3.2 Nests

Food can only be eaten by agents if they are at specific spots in the environment, called nests. The radius of a nest is several times as large as the radius of an agent. Food is usually not positioned in a nest, so to consume a unit of food, agents first need to take it to a nest.

2.3.3 Predators

In some simulations there are also predators present. These predators follow agents they encounter and severely reduce their fitness if they reach them. To keep the amount of random variation in the simulations limited, these predators are hard-coded, they do not evolve. The predators move around randomly and follow agents if they get within a range of 15 squares of them. At each time step, agents that are being followed have a chance to escape from a predator depending on the distance. This chance ranges from 0.1 if the agent is at the maximum vision distance of the predator (15 squares) to 0 if the agent is close by. In some simulations, groups of agents deter predators. This is explained in more detail in Section 2.2.

2.3.4 Obstacles

In some simulations, obstacles is present. These are stationary, impassable objects which prevent the movement through them.

2.3.5 Agents

In the simulation, simple agents attempt to harvest energy and avoid predation. These agents are the object of this study, and their potential swarm behavior is what is investigated in this thesis. The agents are kept deliberately simple so they can not perform higher order tasks on their own. Instead, only emergent group behavior might be able to develop more complicated behavior. The agents do not have any sense of memory and are not able to learn during their lifetimes, their behavior will only change over the generations through the evolutionary mechanics. The agents are able to perform a certain number of actions, which is described in the Section Actions. The agents will all have an energy level, which depletes over time and by performing actions, and is replenished by consuming food. The energy has no direct effect on the agents during a generation, but is important for the survivor selection, which is discussed in Section 2.3.6.

The behavior of the agents is determined by two separate mechanisms, depending on the situation. While an agent is in the same square as a food source and is not currently carrying food, it is hard-coded to pick the food source up, walk with it to a nest, drop it and then eat the food. In all other cases, the behavior of the agents is determined by a neural network,of which the weights are also the genes of the agents. The behavior in some situations is hard-coded because learning the behavioral sequence of moving food to a nest is not part of this research. Instead, the research is about the potential emerging of group and swarm behavior among agents searching for food. Hard-coding the returning of the food to the nest makes it easier for the agents to perform the basic task of the environment, namely gathering and consuming food. Only if the agents can perform this task on their own is group behavior likely to start emerging. Secondly, it also makes the behaviors easier to analyze, because only the task of finding food sources is being evolved. At the start of each generation, the agents are placed at random positions in the environment. The different properties of the agents are discussed below.

Energy

Each agent starts with 500 energy, and consumes energy per iteration. The base energy consumption per iteration, also known as the energy decay rate, is set to 1. The agents consume additional energy by performing certain actions. These actions and their associated energy costs are discussed in more detail in the Section Actions. To replenish energy, the agents are able to consume food they encounter. Each unit of food consumed replenishes the energy by 250. There is no upper or lower cap present for the energy. This is to enable the evolutionary algorithm to distinguish between different gradients of success.

Sensors

The agents each have a number of sensors. The sensor input corresponds directly to the input nodes of the Neural Network. The value of each node is determined by the number of relevant objects detected and the range from the agent to the objects detected. The input value for a sensor is 0.25 for an object if the object is very close and 0.125 if it is at the maximum range, anything in between has a value between 0.25 and 0.125 proportional to the distance from the agent. The total value for a sensor is the sum of the values for all objects it detects.

The agents have 3 different sensors. One of these is the proximity sensor. The proximity sensor has a range of 10 squares, and can detect any type of entity. In the environment, 5 different types of entities can be observed: food, nests, agents, predators and obstacles. The proximity sensor is able to detect 2 different kinds of agents, which are treated as two different kinds of entities for the sensors: agents with an energy level of at least $1/4$ the maximum energy and agents with less than $1/4$ of the maximum energy. This makes for a total of 6 different kinds of detectable objects: agents with low energy, agents with high energy, food, nests and predators. Each of these objects maps to a different input node, so the proximity sensor uses 6 different input nodes. Another sensor is the wall sensor. The purpose of this sensor is to detect whether the agent is at the edge of the field. For each direction, up, down, left and right, the wall sensor has an input node which represents whether the agent is near a wall in that direction. On top of that, the agents also have a sense of smell. The purpose of this is to allow the agents to approach food or stigmergic traces from a relatively large distance from all directions. The smell sensor detects food and stigmergic traces in each direction for up to a distance of 25. For each direction, an input node is reserved for the presence of food and for the presence of stigmergic traces. This makes for a total of $2 * 4 = 8$ olfactory input nodes in the Neural Network. This sensor allows the agents to sense stigmergic traces and food sources from a distance, and to determine in which direction they are.

Internal States

The agents have a number of internal state values which are represented by input nodes in the Neural network. Each agent has an internal state which represents its energy level. The value of the corresponding input node is equal to the energy level divided by the starting energy level of 250. The agents also have a number of binary internal states. These can be either 1 (yes) or 0 (no). These internal states represent whether the agent is currently in the same square as a food source, whether it is in the same square as a nest and whether it is currently carrying food.

Actions

The agents can take the following actions:

- IDLE: The agent does nothing.
- UP: Move to the square above this one, unless an obstacle or wall is present there.
- DOWN: Move to the square below this one, unless an obstacle or wall is present there.
- LEFT: Move to the square to the left of this one, unless an obstacle or wall is present there.
- RIGHT: Move to the square to the right of this one, unless an obstacle or wall is present there.
- TOFOOD: Move in the direction of the closest food source within the proximity range. The agent moves in the dimension (x or y axis) in which it is furthest seperated from the goal.
- TOAGENT: Move in the direction of the closest agent within the proximity range. The agent moves in the dimension (x or y axis) in which it is furthest seperated from the goal.
- TOHIGHFITNESSAGENT: Move in the direction of the closest agent with a fitness of atleast 250 energy within the proximity range. The agent moves in the dimension (x or y axis) in which it is furthest seperated from the goal.
- AWAYFROMPREDATOR: Move in the opposite direction of the closest predator within the proximity range. The agent moves in the dimension (x or y axis) in which it is closest to the goal.
- RETURN: The agent moves towards the nest. The agent moves in the dimension (x or y axis) in which it is furthest seperated from the goal.
- PICKUPFOOD: The agent picks up food.
- DROPFOOD: The agent drops the food it is carrying.
- EATFOOD: The agent consumes food. This replenishes 250 energy.

  If an agent attempts to perform an action but is unable to do so, for example if it tries to move to a certain type of entity while there are none in its proximity range, the agent does nothing that time step.

Stigmergy

Agents automatically leave behind stigmergic traces in the environment. These can be seen as representing bodily waste products, and are important for the research into swarm behavior. Each square has a stigmergic value, which ranges from 0 to 1000 and starts at 0. At each time step, all agents add 1 to the stigmergic value of the squares they are on. Each time step the stigmergic value of each square decays with 0.1 percent.

Neural Network

The behavior of the agents is determined by a neural network. Neural Networks are a common choice for directing the behavior of agents in the field of evolutionary simulations. Nolfi and Floreano [2001, pp.39] cite several reasons for this. One such reason is that Neural Networks offer a relatively smooth search space. "Gradual changes to the parameters defining a neural network (weights, time constants, architecture) will often correspond to gradual changes of the behavior)" Nolfi and Floreano [2001, pp. 39]. Another important point they cite, which is very relevant to this research, is that "Neural Networks can be a biologically plausible metaphor of mechanisms that support adaptive behavior. They are a natural choice for those researchers interested in replicating and understanding biological phenomena from an evolutionary perspective" Nolfi and Floreano [2001, pp. 39].

As to allow space for emergent behavior, the nodes in this Neural Network do not map 1 on 1 with higher-order behavior. No behavior of the agents that is directed by the neural network has been pre-determined. Rather, the nodes represent the input and output mappings of the agents. The input nodes are filled with the values from the sensors and the internal states of the agents. The output nodes correspond to the different actions it can take, minus the actions the agent is hard-coded to sometimes perform. Together, this amounts to 23 input nodes and 10 output nodes. Both the input and the output is discussed in subsequent sections. There is also a single hidden layer present in the network. This hidden layer is added to allow for behavior to emerge which is not linearly dependent on the input. For example, this enables agents to override a certain behavior if it is in front of a pheromone trail, which it will decide to follow instead. Because the behavior of the individual agents should remain fairly simple, only one hidden layer is added. Another reason for only adding 1 hidden layer is that adding extra layers adds unnecessary complexity, making it harder to analyze. Given the simplicity of the environment, a single hidden layer should be powerful enough for this research. In practice, few problems that cannot be solved with 1 hidden layer can be solved with more. The amount of nodes in the hidden layer is equal to $2/3$ the amount of input nodes + the amount of output nodes[Heaton, 2008], which makes for 25 hidden nodes. To limit the size of the search space, the weights of the network are encoded as integers in the interval $[-100 : 100]$. The weights are initialized at random values in the interval $[-5 : 5]$.

If the previous action of the agent was to move up, down, left or right two adjustments are applied to the activation values of the input nodes after they have been computed by the neural network. Firstly, the activation value of the opposite move action of the previous action is reduced by two times the absolute value of the current activation value that action. The actions UP and DOWN are considered opposites to each other, as well as LEFT and RIGHT. This is to prevent the agents from continuously walking in the same place. A second purpose of this is to make it easier for the agents to follow stigmergic paths. If an agent is halfway through following a long path, there is as much stigmergic activation behind it as ahead of it. This would make the agent as likely to go back as to go forwards, which is not intended. The second adjustment that is performed on the activation values of the output nodes if the agents last action was to move is to add the absolute value of its current activation value to the last move action that was performed. This is only done 5 times in a row, if the agent keeps performing the same move actions after that its activation value will not be increased. If a new move action is performed after that its activation value is increased up to five times again. This adjustment is performed to encourage the agents to move in the same direction repeatedly, which makes it more likely the agent will cover larger distances and find food sources. An additional purpose of this is to make it easier for agents to follow a curved stigmergic trail, as and agent is more likely to pursue the new direction after it has changed direction.

The activation values of the hidden nodes are determined by the sigmoid function applied on the normalized summed input from the input nodes and input-hidden weights and the bias weights. The sigmoid function is a well known method for mapping a value to the $[0 : 1]$ interval, and is defined by the following equation:

$$S(t) = \frac{1}{1 + e^{-t}} \tag{2.1}$$

. The input in the hidden layers is normalized to prevent over-saturation of the sigmoid function. At each time step, when all hidden nodes have summed their inputs from the input nodes and the bias, the total list of all hidden node activations is normalized to have an average deviation of 1 from 0. After that, the sigmoid function is applied.

The actions that are performed are chosen rank based and stochastically. When an agent chooses an action, the activation values of the output nodes of its neural network are compared to each other, and the action which corresponds to the node with the highest value has a predetermined chance of 0.9 to be chosen. If it is not chosen, the action which corresponds to the node with the second highest value will have a chance of 0.9 to be chosen. This process can potentially repeat itself until there is only one action left, which is then automatically chosen.

More formally: Let A be the set of all possible actions, which is also the set of all output nodes in the neural network, and $r(a)$ be the rank of an action $a \in A$ in the interval $[1 : |A|]$, where 1 means it has the highest activation value. The probability $P(a)$ that a is chosen is calculated with the following formula:

$$P(a) = \begin{cases} (1-0.9)^{r(a)-1} & \text{if } r(a) = |A| \\ 0.9 * (1-0.9)^{r(a)-1} & \text{otherwise} \end{cases} \tag{2.2}$$

Note that this formula satisfies the constraint that the probabilities for all actions sum to 1, or

$$\sum_{a \in A} P(a) = 1 \tag{2.3}$$

2.3.6 The Evolutionary Dynamics

There is a preset number of generations. Each generation, the best agent is kept, and the rest of the new generation is filled in with offspring. Agents are never removed during a generation, the agents that are present at the start continue to operate for the entire generation.

The fitness of the agents is determined by their energy level and on whether it is caught by a predator that simulation run. A higher energy level will linearly increase the fitness. If there are predators in the environment, the fitness values of agents that have been caught are decreased by 0.5. The formula for the fitness is the following: $U = E/500 - 0.5 * P$, where U is the fitness, E is the energy level at the end of a generation and P is a variable which represents whether the agent was caught by a predator. P is 1 if the agent was caught by a predator, and 0 otherwise. The fitness function has deliberately been made simple as to give room to different kinds of emergent behavior. “The more detailed and constrained a fitness function is, the closer artificial evolution becomes to a supervised learning technique and less space is left to emergence and autonomy of the evolving system” [Nolfi and Floreano, 2001].

The parents are chosen by a rank based method. The agents are ranked according to their fitness and each agent is given a ranking score. The ranking score of an agent is equal to the total number of agents minus the number of ranks this agent is removed from the best agent, so the the best agent has a ranking score equal to the number of agents, second best one equal to the number of agents minus 1, and the worst has a ranking score of 1. For each parent that is selected, each agent has a chance to be selected equal to its ranking score divided by the total ranking scores of all agents. It should be noted that the same agent can be selected multiple times.

Apart from the agent that is preserved, the new generation is filled by copying the selected parents directly into it, and mutating their genes. The mutation is adding a random value from a normally distributed function with a mean of 0 and a standard deviation of 0.05. There is no crossover operator present. This is for the reason that crossover may generate instability and can lower performance. ”Over the last years, much more attention is being paid to the mutation operator.”Nolfi and Floreano [2001]. The mutation operator is applied separately on all the genes of the agents.

2.3.7 The standard values

In this Section, the default settings are displayed. Unless indicated differently, these values is applied in the simulations.

Table 2.3: General Settings

Setting	Value
FieldHeight	100
FieldWidth	100
NrAgents	25
NrObstacles	0
NrNests	1
NrPredators	5
maxFood	50
PredatorSensorRange	15
NrActionsPerGeneration	500
NrGenerations	500

Table 2.4: Agent Settings

Setting:	Value
NrSensorValues/InputNodes	23
NrHiddenNodes	25
NrActionTypes/OutputNodes	10
InitalWeightSize	[-5:5]
MaxWeightSize	[-100:100]
Starting Energy	500
EnergyReplenishmentFood	250
Energy Decay	1
Proximity Sensor Range	10
Wall Sensor Range	10
olfactory Range	20

2.4 Predictions

In this Section, I will first return to the theories of the emergence of group and swarm behavior in general and place them in the context of the specific environment and mechanics used in this research. Afterwards, I will discuss the 3 main behaviors that are being researched (Clustering, Following and Stigmergy). For each of these behaviors, I will discuss if, under what circumstances, and how they may emerge.

2.4.1 Groups

In this subsection, I will theorize about the mechanisms that could cause the emergence of groups in the simulation environment, and how they could influence the behavior of those groups. The mechanisms for regular group forming that were discussed earlier are the following:

- Kin Selection: This seems unlikely to be a cause for group behavior in this simulation, because there is no explicit relatedness parameter present. Besides that, since all individuals are initialized at random spots in the environment each simulation, individuals who happen to have similar genes will be just as likely to form groups with very different individuals as with each other.

- Direct/Indirect Reciprocity: this will not play a role in this simulation simply because the agents lack the sense of memory and cognitive capabilities required for it.

- Spatial Reciprocity: this could be a cause for group behavior in this simulation. If multiple individuals depend on each other for survival, as they very well could in this simulation if the Predator en Group Protection parameters are active, it would be in their interest to help each other. However, "Helping" is a rather vague concept in the terms of this simulation, as the only way for agents to protect against predators is to cluster, which protects all agents in the cluster.

- Group Selection: This could also be a cause for group behavior in this simulation. If individuals who group together perform better than other individuals, they would be favored by the evolutionary mechanics of this simulation.

- Profiting: If individuals can profit from the behavior of other individuals, groups can be formed. This can be present in this simulation, for example by agents following other agents towards food sources.

2.4.2 Swarms

The above explanations are merely about group behavior. For swarm behavior to emerge, more thorough cooperation is necessary. Swarm behavior is the collective behavior of decentralized individuals, not merely individuals who help cooperate. A higher order behavior has to emerge. The theorizing I did in Section 1.3 should apply to the simulation. I will now provide a quick recap of said theorizing.

If individually operating (non-swarm) individuals perform actions in certain steps, which change the environment, they could use their own previous actions as clues for what to do next. If another individual meets the environment in an altered state, it can just continue the action from the step the process is in. A problematic feature of some swarms is that individuals do not function on their own. This can be evolved by means of intermediate steps in the evolution of the group. After an intermediate step, an individual can count on a certain behavior of the others and therefore change its behavior into something that would not work in isolation.

I expect that stigmergic trace following behavior will emerge this way in the simulations where stigmergy is active. In these simulations, the agents automatically leave behind stigmergic traces, which represent bodily waste products and are thus consistent with the exaptation theory if they are utilized for a following purpose later. If the food is clustered and the agents reach these clusters relatively often, stigmergic trails should lead from the nests to the food clusters. If this is the case, the agents can potentially learn that to follow the stigmergic traces from the nest will result in finding food quicker.

- I expect that having predators and group protection active will make agents more likely to cluster.

- I expect that if food is clustered, the agents will be more likely to perform following, clustering or stigmergic following behavior, because entire groups can feed at the same general area. Otherwise, groups would deplete local food sources.

- I expect that the ability to detect stigmergic traces in the environment will make agents follow them if food clustering is active.

2.4.3 Clustering

In this subsection, I will provide a detailed account of how Clustering behavior could emerge in the simulations. This will be more detailed than the predictions about the other behaviors, because many observations I will make here with regard to the neural network and the evolution will also apply for following and stigmergic behavior.

Clustering is a specific kind of group behavior. It is not swarm behavior, but it could later cause swarm behavior to emerge. The fact that the agents can have an interest in each others survival, to keep the group alive, means that they could later evolve a more thorough cooperation through mechanics like kin selection and spatial reciprocity.

Parameters

The parameters that I expect to be required for clusters to emerge are the following: Predators, Group Protection and Food Clustering. Although there are many possible ways in which groups can emerge, some of which were described in the subsection Groups, most of these seem unlikely to occur in a "neutral" environment such as the one in this thesis. After all, there are many solitary creatures in nature, none of the possibilities for evolving group behavior are in any way inevitable. This is why this research focuses on the Group Protection theory for grouping. If there are predators in the environment and agents are less likely to be attacked by them if they are in close proximity to each other, they have a concrete reason to band together. Once they group, other group forming mechanisms like group or kin selection could occur.

I expect the Food Clustering Parameter to be beneficial for group forming because this will reduce the problem of local depletion of food sources. If the food is scattered evenly in the environment, groups will quickly deplete local food sources, possibly before feeding all members. In this case, it may be more efficient for the members to spread out, resulting in the opposite effect as group forming. Depending on the size of the groups of agents and the clusters of food, having food clustering active may negate this issue.

Neural Network

The theory behind the emergence of groups in this research is that the parameters Predators, Group Protection and Food Clustering will cause the agents to form groups. On a more abstract level, this means the natural selection has to favor agents whose neural network weights cause them to respond to the environment in such a way that they tend to move towards other agents. The agents start each generation with randomized neural network weights, which are also their genes. Some agents have genes which will cause them to move towards the other agents, some do not. Each agent has an input node for all other object types in each position. For an agent's genes to cause it to move towards other agents, a positive connection between detecting an agent or agents with the proximity sensor and performing the "Move To Agent" action is required. This means there have to be positive connections between the sensory node that detects another agent to hidden nodes which have positive connections with the "Move to Agent" output node.

Evolution

If the agents that move towards each other as described above tend to have a higher fitness than the agents that do not, they will be more likely to be selected as parents for the next generation. Since there will be no crossover operator employed in this simulation, the selected networks will be mutated and then passed on to the next generation. This means that the relevant genes for moving towards other agents will thus be extra represented in the next generations, resulting in networks that are more likely to cause agents to move towards each other. The crucial point here is that agents that tend to move towards other agents must have a higher fitness than those that do not. With the predators and group protection parameters on, agents that group together are less likely to be followed by predators. Being caught by a predator greatly reduces the fitness of an agent, so clustering would increase the utility.

However, since there is no one on one mapping between genes and behaviors, there will inevitably be side effects to the changing of the genes in the neural network. For this reason, it is uncertain if genes which cause agents to move towards each other give a "net" benefit to the utility. For example, if agents have a strong inclination to move towards each other they may decline to move to valuable food sources, instead opting to move to other agents. With this behavior, the utility of these agents would decrease, thus making them less likely to be selected for reproduction. A more optimal behavior would be to only move towards successful other agents if they can not detect any food sources. This behavior is possible for the neural network to develop through the hidden layer. The presence of both an agent and a food source could result in the hidden layer overriding the input from the agent and providing a positive connection to the output nodes for moving towards the food source. However, this behavior is far from guaranteed. If the behavior to move to other agents will work detrimental at first, since the aforementioned more advanced behavior has not been evolved yet, the agents might not learn to move to other agents. In other words, the agents could get stuck in a local minimum. With the parameters in place, a variety of behaviors and behavioral conflicts may arise.

2.4.4 Following

Like clustering, following behavior in general is not actual swarm behavior, but may lead to more thorough cooperation later.

Parameters

I expect the Food Clustering Parameter to be necessary for following behavior to emerge. If the food is scattered around the world, any food that is found by an agent will be unlikely to be close to another. If the agent that is being followed will take the food itself, following and agent will actually be detrimental to the chances of finding food. If the food is clustered, following agents with a high energy value may be positive for the chances of finding food, since their high energy value indicates these agents are good at finding food sources.

Neural Network

The prediction behind the emergence of following behavior in this thesis is that agents will evolve a behavior of following successful agents if food clustering is active. On a more abstract level, this means the natural selection has to favor agents whose neural network weights cause them to move towards agents who have a high energy level. Since following behavior is achieved by the action "Move To Successful Agent", an indirect positive connection between the input node of observing a successful agent and the action "Move To Successful Agent" has to be present.

Evolution

If the agents that move towards successful agents have a higher fitness value than the agents that do not, these agents will be favored by the evolution. I expect this will be the case because successful agents will generally be better at finding food sources than other agents. One weakness of following behavior is that if all agents will perform it, there will be no agents left to locate the food sources, resulting in a low performance by all agents. Following behavior is only effective if some agents perform it and some will explore the environment. This division could emerge, because agents have a separate action for moving towards a successful agent and moving towards another agent. The agents could potentially explore the environment unless a successful agent is within vision, which they will then follow. Whether this behavior will result in a higher fitness on average remains to be seen.

2.4.5 Stigmergy

Stigmergic following behavior is the only behavior that is investigated in this research that can be considered to be true swarm behavior.

Parameters

I expect the Food Clustering Parameter to be necessary for stigmergic following behavior to emerge. If food clustering is active, there will be areas in the world where several food sources are present. If some agents have found this food and brought it back to the nest,the squares between the nests and the food clusters will have stronger stigmergic values than the other squares. This makes it beneficial for agents to follow the trails.

Neural Network

For stigmergic following behavior to be displayed, agents will have to move in the direction where they perceive a high stigmergic value. This means the natural selection has to favor agents who perform the action to move in one of the four directions if the smell sensor detects a strong total stigmergic value in that direction. This behavior will emerge if there are indirect positive connections between the input nodes in the neural network which represent the stigmergic activation values in the different directions and the output nodes which represent the actions to move in those directions. The fact that agents will be unlikely to perform the move action that is the opposite of the move action it has just performed (see Section 2.3.5), means it will not be likely for them to reverse their direction on the trail, which could enable them to follow the trails to the food clusters.

Evolution

A requirement for stigmergic following behavior to be beneficial for the agents, is that the strongest stigmergic traces lie between the nests and the food clusters. Since agents will directly move to the nests once they have picked up food, the squares between the food clusters and the nests will be passed by agents more often than other squares. Because agents always drop stigmergic traces on the squares they are on, the squares between the nests and the food clusters should have higher activation values. Learning to follow these trails reliably to food clusters would give agents an evolutionary advantage, as this would negate the need to search for food if other agents have found a food cluster. Following weaker stigmergic trails, on the other hand, would likely not result in a high fitness value. Since agents automatically leave behind stigmergic traces, most squares will have some stigmergic value. This includes squares which are not between a nest and a food cluster. To successfully utilize stigmergic traces for foraging behavior, agents should only follow traces of a certain minimum strength, and display normal searching behavior otherwise. Since the total strength of the stigmergic traces in a direction directly determines the strength of the corresponding input node in the neural network, this behavior can be achieved by the proper weights in the neural network.

Chapter 3

Results

All experiments were run with 10 tests per genotype with 10 genotypes per setting, which means $n = 100$. Three experiments were run, one to investigate following and clustering behavior, one to investigate stigmergic behavior, and one extended test to test for a possible decrease in fitness in some simulations (see Section 3.3). For each experiment, the baseline test simulations were compared with the experimental test simulations. The baseline condition is the condition in which the agents are unable to perceive each other and stigmergic traces (see Section 2.1.2). For the first experiment, three right-tailed t-tests were performed for each setting that was investigated: one with the clustering as dependent variable, one with the following, and one with the fitness. For the second experiment, two right-tailed t-tests were performed for each investigated setting: one with the stigmergy as dependent variable and one with the fitness. This means three t-tests were performed on the twelve settings in the first experiment, and two t-tests were performed on the three settings in the second experiment, which makes for a total of 42 tests. Because 42 tests were run, the significance level α is set to $0.05/42 = .00119$. For a setting to yield a significant result, it is also required for t to be greater than 0, since these are right-tailed t-tests in which a certain behavior is said to have emerged if the experimental condition has a higher operationalisation value than the baseline condition.

3.1 Clustering and Following

For the first experiment, three parameters were varied (see Table 2.1). The different simulations will be referred to by a three number code, which represent the values of the parameters, with the first value referring to the Predators Parameter, the second value to the Group Protection parameter and the third value to the Food Clustering Parameter. A binary variable was represented by 1 if it is active and 0 if it is not. See Section 2.2 for the meanings of the different values of the parameters. The average fitness for a single test run is considered a single value for the purpose of the experiment. The average fitness per experiment is the average of the average finesses for the different test runs for that experiment. This is done in the same way as with each setting having a single operationalisation value for each operationalisation method (clustering, following and stigmergy). The fitness averages over all parameter settings during the training simulations are shown in Figure 3.1. Note that these fitness values are taken from the training simulations, since test simulations were only run after all train simulations were completed. The means of the operationalisation values and the fitness for the first experiments are shown in Figures 3.2, 3.3 and 3.4. These values are the averages over all settings in the first experiment during the evolution. The results of the statistical tests are shown in Table 3.1. No significant results were found.

Figure 3.1: Average Fitness over generations Experiment 1.

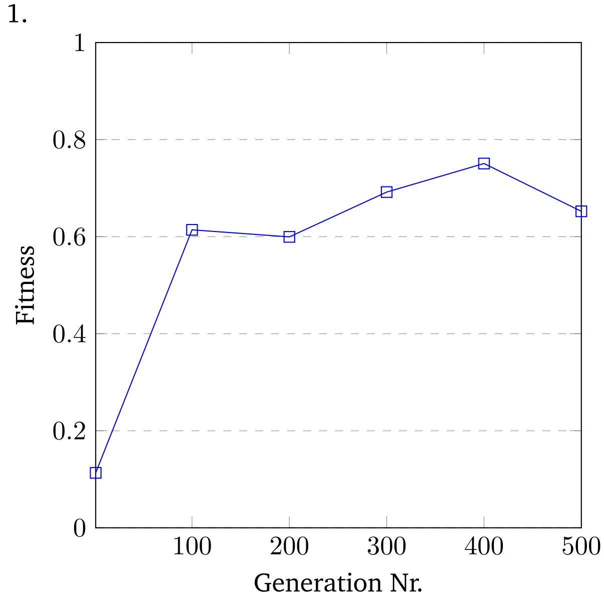

Figure 3.2: Clustering Experiment 1. The experimental settings are set on the x-axis, and the average Clustering value per test run of that experiment are set on the y-axis

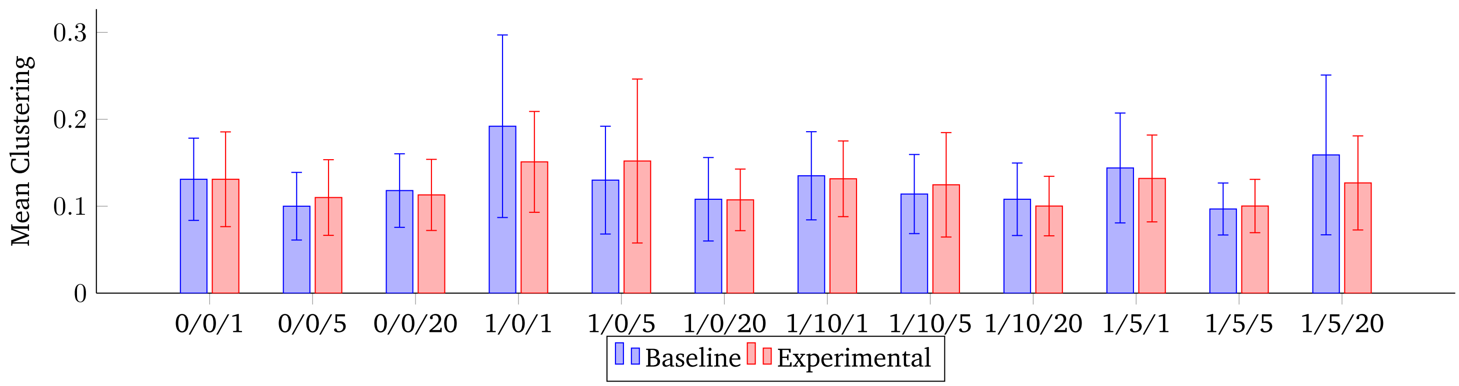

Figure 3.3: Following Experiment 1. The experimental settings are set on the x-axis, and the average following value per test run of that experiment are set on the y-axis

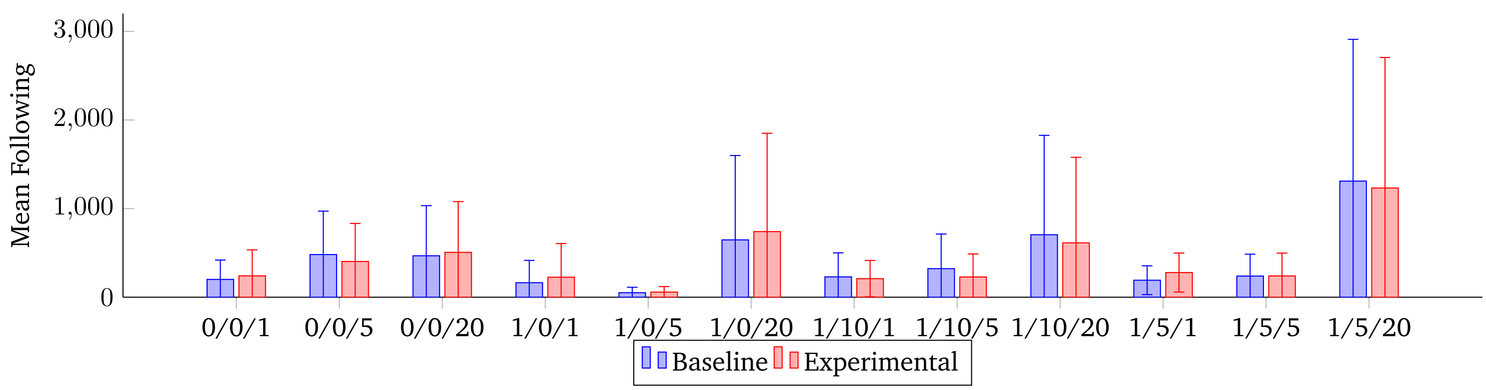

Figure 3.4: Fitness Experiment 1. The experimental settings are set on the x-axis, and the average Fitness value per test run of that experiment are set on the y-axis

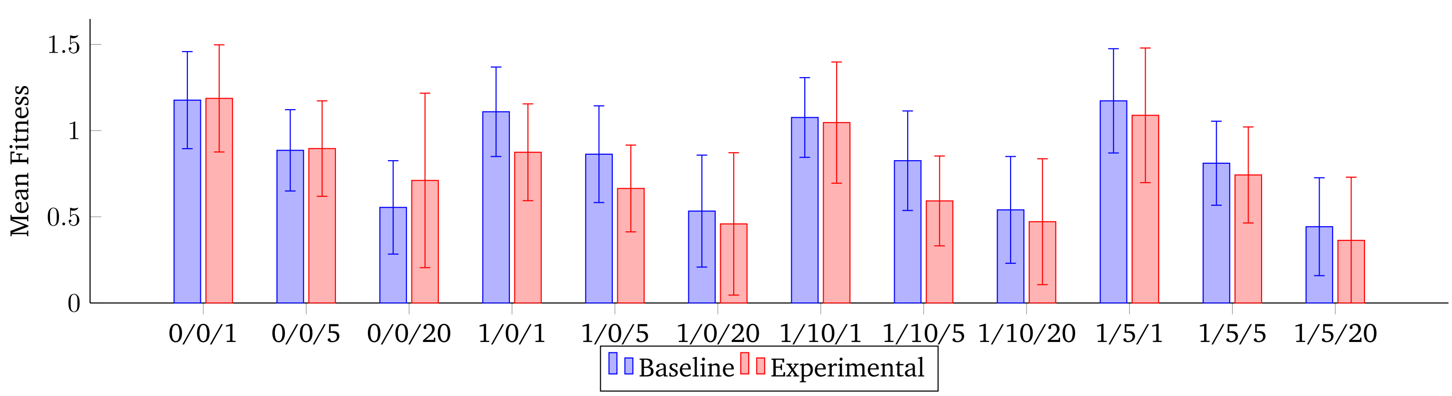

Figure 3.5: Average Stigmergy Experiment 2. The setting number represents the food cluster size in the environment.

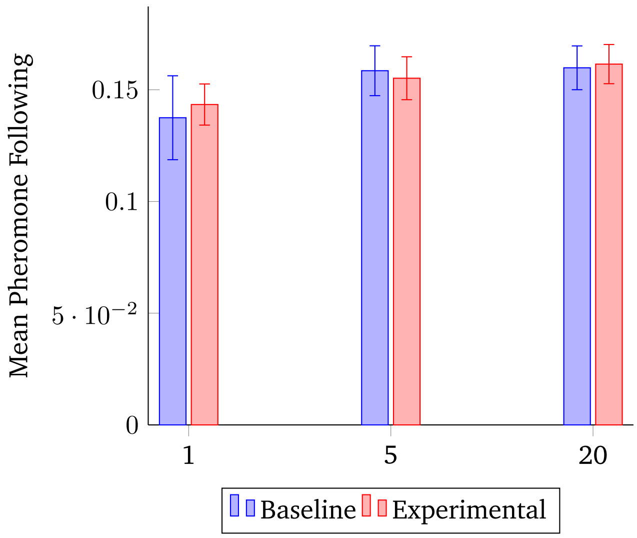

Figure 3.6: Average Fitness Experiment 2. The setting number represents the food cluster size in the environment.

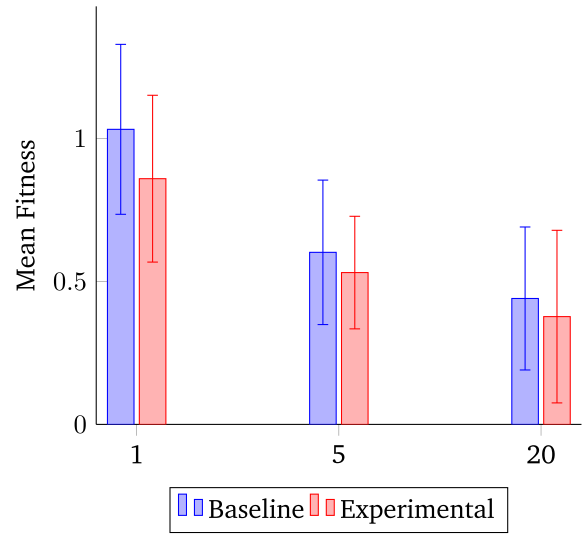

Table 3.2: Test Results Experiment 2

Cluster Size	pStigmergy	tStigmergy	pFitness	tFitness
1	0.0114	2.30	0.00179	-2.95
5	0.0389	-1.77	0.0379	-1.79
20	0.165	0.98	0.137	-1.10

3.2 Stigmergy

For the second experiment, to investigate the stigmergic behavior, standard values were used for most parameters (see Section 2.1). Only the Food clustering Variable was varied, the others were at the default for experiment 2 as defined in Section 2.1. Each parameter setting in this experiment is defined by the value of the Food Clustering Parameter, which will also be used to refer to a setting. Statistical tests were performed to investigate a potential increase in Stigmergic activity and fitness. The results of the statistical tests can be found in table 3.2. No significant effects were found. The means of the operationalisation values and the fitness values can be found in Figures 3.5 and 3.6. The average fitness values of the training data over the generations are shown in Figure 3.7.

Figure 3.7: Average Fitness over generations Experiment 2.

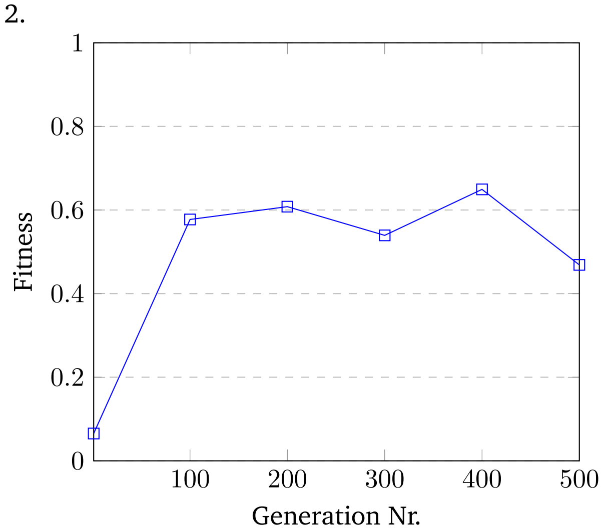

3.3 Extended Research

One unexpected observation in the results for a number of settings in both experiments is the fact that the p-value for the fitness is low, but the t-value is negative. This is unexpected because this would mean the agents perform better if they are not able to observe each

Table 3.1: Statistical Results Experiment 1

Setting	pClustering	tClustering	pFollowing	tFollowing	pFitness	tFitness
0/0/1	0.469	0.0774	0.212	0.800	0.423	0.188
0/0/5	0.138	1.09	0.157	-1.01	0.415	0.216
0/0/20	0.266	-0.625	0.361	0.355	0.0408	1.75
1/0/1	0.0127	-2.25	0.218	0.782	<0.0001	-4.931
1/0/5	0.155	1.016	0.285	0.570	<0.0001	-4.30
1/0/20	0.454	-0.117	0.326	0.452	0.184	-0.904
1/10/1	0.363	-0.352	0.337	-0.422	0.299	-0.528
1/10/5	0.210	0.809	0.0614	-1.55	<0.0001	-4.29
1/10/20	0.146	-1.056	0.358	-0.364	0.149	-1.044
1/5/1	0.151	-1.03	0.00581	2.547	0.115	-1.204
1/5/5	0.274	0.603	0.489	0.0285	0.0820	-1.40
1/5/20	0.0310	-1.88	0.378	-0.312	0.110	-1.23

Table 3.3: Test Results Extended Experiment

Setting	pFitness	tFitness
1/0/1	<0.000001	-15.6
1/0/5	<0.000001	-10.8
1/10/5	<0.000001	-13.9
S1	0.000138	-3.69

other. Since the t-tests that were performed were one-sided and only tested whether the fitness value would be higher in the experimental condition than in the baseline condition, it is not possible to tell from this data if the average fitness is really lower in the experimental condition than in the baseline condition in those settings. For this reason, a new experiment was run to test for this. For the 4 settings which have a p-value that would have been significant if the t-tests where left tailed, new simulations and tests where run. The fitness values from the resulting data were analyzed with a right tailed t-test for each setting. Because this is a test to validate earlier results which only tests a limited number of settings, the number of simulations per setting was increased from 10 to 20 and the number of tests per simulation was increased from 10 to 50. The significance level for these tests is set to $0.05/(42+4) = .00109$. The Statistical Results are shown in Table 3.3. The average fitness values from the Baseline and Experimental condition are shown in Figure 3.8, and the average fitness values of the train data for all settings over the generations are shown in Figure 3.9. The binary codes of 3 numbers are settings without stigmergy, encoded in the same way as in Experiment one, and the setting S1 means Stigmergy is active, food clustering is 1 and the other two parameters are at the default for Stigmergy. A significant effect was found for all four tests. This means that the fitnesses of the agents are greater in the baseline condition than in the experimental condition.

Figure 3.8: Average Fitnesses Extended Experiment

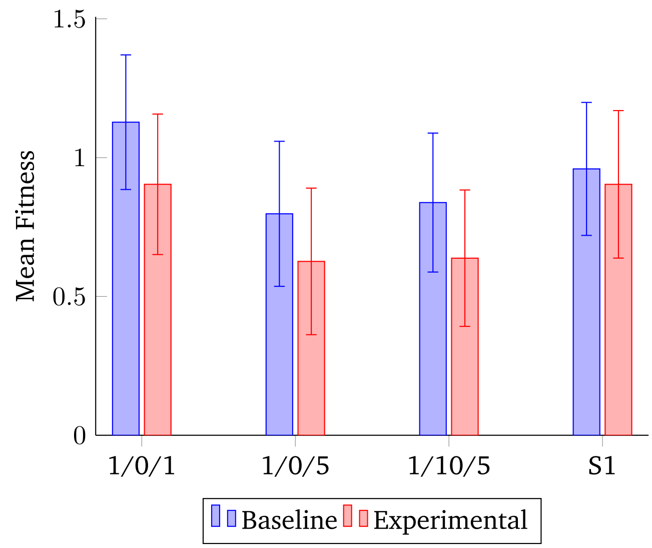

Figure 3.9: Average Fitness over generations Experiment 3.

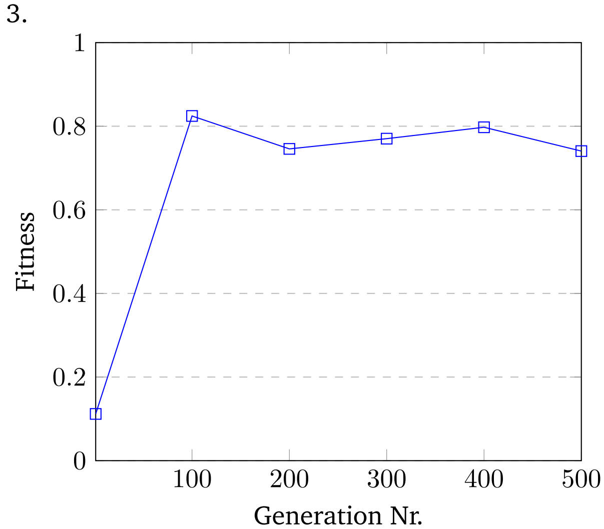

Chapter 4

Discussion

4.1 Overview

Swarm Behavior is the collective behavior of decentralized, self-organized systems, natural or artificialZhang et al. [2013]. An interesting question about Swarm Behavior is how, and in what circumstances, it can emerge. There are several potential ways in which swarms can evolve: from individuals or from groups. Several kinds of group and swarm behavior have been identified: clustering behavior, following behavior and stigmergic following behavior. The aim of this project was to investigate by means of a computer simulation in a controlled environment under which of a number of environmental circumstances these behaviors will be able to evolve. If agents would develop clustering or following behavior in certain environmental circumstances, the potential evolution of stigmergic behavior from these agents would have been investigated.

No significant results were found in any setting for any of the three researched behaviors or for the fitness values. The Standard deviations are high in comparison to the mean values, especially for the following behavior. Some p-values suggest there is an effect, but with 10 simulations and 10 tests per simulation, which amount to 100 data points, the effects are not significant. The first three Research Questions were:

- **Research Question 1:** with which parameter settings will clustering behavior emerge from novel agents?
- **Hypothesis 1:** with food clustering, predators and group protection active. I expect predators and group protection to be necessary to provide an incentive to form groups. I expect food clustering to be necessary to prevent groups from depleting their food sources too quickly.
- **Research Question 2:** With which parameter settings will following behavior emerge in the simulation from novel agents?
- **Hypothesis 2:** with food clustering active. I expect the agents will follow other agents to find clusters of food.
- **Research Question 3:** With which parameter settings will stigmergic behavior emerge in the simulation from novel agents?
- **Hypothesis 3:** with food clustering active. I expect the novel agents will first learn to follow and later display stigmergic behavior. Once agents have already learned to follow other agents, the following of stigmergic cues can provide a more reliable mechanism to achieve the same goal, the finding of food clusters.

These three research questions are answered with: no significant effects were found. The other Research Questions were optional and have not been investigated, since they required a significant effect for at least one setting for a previous research question. See Section 1.1.

4.2 Lack of Significance

There are several possible explanations for the lack of significant results. The simplest explanation is that with the environment, agents, parameters and evolutionary mechanics in this research the agents do not develop clustering, following or stigmergic following behavior. One possible reason for this is that the agents may get stuck in local maxima, as was mentioned in Section 2.4.3. Since agents are selected purely on their performance each generation, the evolutionary mechanics do not encourage the evolution of behaviors which are not directly beneficial. This could mean that agents who move towards other agents will generally perform worse than agents who do not because they occasionally prefer moving to other agents over moving to food sources. This may cause the agents to never develop the more optional behavior to move to other agents if they do not detect a food source. See Section 2.4.3 for a more detailed explanation.

Another potential reason for the lack of significant results is that the amount of nodes and connections in the neural network is too high, and the different types of group behavior have too little effect in comparison to the basic foraging behavior. In Sections **??** and 2.4 I have theorized that the behaviors could emerge provided that they have a positive effect on the fitness. Having a large neural network with many nodes increases the search space, making it harder for the evolutionary algorithm to converge to an ideal setting. This makes it even harder to surpass a local maximum, since the right mutation out of a very large number of mutations has to occur. In such a large network, the input to each hidden or output node will be the sum of the output of many input or hidden nodes, each multiplied by correspond weight (see Section 2.3.5). There will inevitable be a large amount of noise in this input. For a certain node to make a difference in the behavior of the agent, it must have a large weight to a certain node and or a large negative weight to the other nodes, and the highest activation values of the nodes in the next layer must be similar to be able to make a difference. For a direct behavior which increases the fitness by a large amount, such as moving towards perceived food sources, this is likely to evolve, the weights can quickly evolve to sufficiently high values to make the neural network usually exhibit the right behavior. In the case of a behavior that is only preferable some of the time or does not increase the fitness by a large amount, such as the move towards successful agent behavior which is pivotal for group behavior, it is harder for a single node to make the difference. This makes it less likely to evolve the group behaviors investigated in this thesis. It is still possible for an evolutionary algorithm to overcome this, but it is possible that the size of the neural network has prevented them from evolving. One way for a future research to find out if this is really the case is to reduce the amount of inputs and outputs into the network and the size of the hidden layer, effectively pruning the network. If the agents are able to develop more sophisticated behavior then, it is likely they were held back in this research by the size of the neural network. Pruning inputs and actions is a delicate matter, since it will also limit the ways in which the agents can directly respond to the environment. I would suggest to remove the predators and nest input cameras, and to remove the move away from predator output node. This would mean that the predators and group protection parameters can no longer be investigated, but it would limit the size of the neural network. Another option is to remove all actions which cause the agent to directly move towards the nearest instance of a certain object (see Section 2.3.5). This would greatly reduce the size of the neural network, but would make it harder for the agents to move towards objects that are within their vision. This will make the basic foraging behavior harder, since the agents will have to learn to move towards objects without these actions, but would reduce the size of the neural network. Some adaptations in the environment may be necessary for the agents to successfully operate without these actions. The operationalisation for the following behavior would need to be redefined, since it is dependent on the "move to agent" and "move to successful agent" actions (see Section 2.1.7).

Another reason for the lack of significance could be the abstraction level. Firstly, any number of features in animals in the real world could be important for the emergence of group and swarm behavior. The aim of this research was to investigate the emergence of these behaviors from simple agents. As a result, the agents are directed by a very basic neural network, and do not have any sense of memory. Individually, the agents do not have much potential to develop sophisticated behavior, given their simplicity and the limited ways of interacting with the environment. This was intentional, since an important feature of swarm behavior in particular is that agents can perform more complicated behavior as a group than they can perform individually. It is possible, however, that the agents are too simple for the investigated behavior to emerge. There could be common features in even the simplest of cooperating animals that were not present in this research. In the Section 2.4 I gave examples of how group behavior may emerge in my simulation, but it may be held back by specific features in my environment. Furthermore, despite the environmental parameters that were investigated, the experiments were conducted in a specific environment. It is possible that certain environmental features that were absent in the simulations play an important role in the development of group behavior. This would mean that the circumstances that were investigated in this research are not enough for the investigated behaviors to emerge in the environment that was used. In Section 2.4 I have predicted that the setup of the agents and the implemented environment are enough to let clustering, following and stigmergic following behavior emerge with some parameter settings, but it is possible that some critical aspects for the evolution of these behaviors were missed in the analysis. This could mean more or different parameters should also be investigated. Future Research could find out if this is the case by replicating my research in an entirely new simulation environment. Every simulation environment is inevitably an abstraction from the real world, but investigating the same settings with the same general parameters with several very different environments would provide insights as to what specific features have an impact on the result. Future research could also investigate new parameters that were not included in this research, to find out whether they have an impact on the outcome.

A different explanation that no significant effects were found is that this research is exploratory and a large number of statistical tests were run. 3 kinds of group behavior were investigated, with 3 parameters that were varied, resulting in 42 tests. A number of the p-values that were found suggest an effect. However, because of the number of tests that were run, the significance level α is so small none of the tests yielded a significant result.

4.3 Observations

In this Section, a number of observations that were made on the data will be discussed. Since no statistical tests were performed for these features of the data that were observed, they may have emerged by chance and not be generalisable to further experiments. However, discussing these features in the context of the theorizing done in this research may still provide valuable insight into the properties of group behavior, and may direct future research.

4.3.1 Fast Convergence

The average fitnesses over the settings in the different experiments rise rapidly from generation 1 to 100 and then fluctuate in both directions, without a clear increase in value. This can be seen in Figures 3.1, 3.7 and 3.9. This means the behavior has likely converged somewhere between generation 1 and 100, so before $1/5$ of the total number of generations that were run. If the behavior converged that quickly, it is likely that the agents have not developed a higher-order group behavior, as that could require more generations to develop.

4.3.2 Lower Fitnesses in the experimental condition

One curious observation of the experimental data is that in most settings, the fitness appears to be lower in the experimental condition than in the baseline condition. For the settings where the p-value was also low and extended research that was done which confirmed this. This is unexpected because the only difference between the baseline and the experimental condition is that the agents are unable to perceive both each other and stigmergic traces. This means that in some settings, the agents with less information perform better. This could be explained by taking into account that the agents have fewer potential active input nodes in the baseline condition. If the agents are unable to perceive certain objects in the environment, some of their input nodes in the neural network are always set to 0. In effect, this decreases the size of the neural network. If the agents are unable to effectively utilise the information they get in the experimental condition but not in the baseline condition, this information can be regarded as noise. This may detract the agents from performing a successful behavior.

If the removal of a few input nodes can lead to an increase in fitness, this suggests that the size of the neural network is indeed detrimental to the exhibiting of an optimal behavior, as speculated about in Section 4.2. This observation also makes it more likely there are no advanced behaviors in the settings where the fitness is lower in the experimental condition, since these are unlikely to evolve if the fitness is lower if they exhibit them, since they are more complicated and would take longer to evolve than the basic foraging behavior. It is possible the agents get stuck in a local maximum while displaying an ineffective group behavior, but this is not likely if this behavior is less effective than the agents acting individually. For future research, I would recommend pruning the neural network, as discussed in Section 4.2.

4.3.3 Fitness and Food Clustering

Another observation is that the general trend appears to be that the fitness decreases if the food clustering increases. See Figure 3.4 for an overview of the average fitnesses in different parameter settings. This may be the case because agents are unable to find isolated clusters, thereby missing them entirely. If the agents distribute themselves over the environment, it will be beneficial for the average fitness if the food is also distributed over the environment, so different agents can find different food sources. Since the amount of food sources that are in a cluster does not increase the range at which they can be detected, a food cluster will not be much easier to find than a single unit of food. Since the same total number of food sources present in the environment is always the same regardless of the food clustering parameter, there will be less clusters if the food clustering parameter is at a high value than there will be if it is at a low value. This means that if the food is strongly clustered, many agents will not find any food sources and some agents will find large clusters of food. Since an agent can only transport a single unit of food to a nest, finding a single unit of food is as valuable as finding a food cluster. This could explain for the apparent drop in fitness averages with an increase in the food cluster size. These effects could be counteracted by group behavior such as following of other agents and stigmergic traces. If the agents are able to follow each other to food clusters, or to revisit food clusters by following stigmergic traces, the fitness would be expected to be higher. However, it appears these behaviors have not evolved, or at the very least not enough to counteract the detrimental effect of food clusters on the fitness values.

4.3.4 Following and Food Clustering

Unlike the fitness, it appears there is a positive connection between the size of the food clusters and the performed following (see Table 3.1). This is in line with the Prediction about following in Section 2.4.4, which states that Food Clustering is necessary for following behavior to emerge because it will only be useful to follow agents if they are heading for clusters of food. If the food is not clustered, following agents will often be detrimental since the agents that are being followed may take the only food in the vicinity for themselves, leaving the following agents in a barren area.

4.3.5 Stigmergy

Although no significant effects were found, the p-values for stigmergic behavior in Table 3.2 are generally lower than those for following and clustering behavior in 3.1. Since these p-values were still insignificant, this does not necessarily mean that stigmergic following behavior emerged, but it could mean that stigmergy has emerged while the other behaviors have not. This would be contrary to expectations, as stigmergy was seen as the most advanced and difficult to evolve type of group behavior in this research. If it is true that Stigmergy has evolved and the other behaviors have not, this would mean the hypothesis that stigmergy would evolve from following behavior, as predicted in Section 1.1.1 is incorrect. In Section 1.3, the possibilities of swarm behavior emerging from "intermediate" group behavior or directly from individuals were discussed. This observation makes it likely that stigmergic behavior can emerge from individuals. Two reservations have to be made about this. It is always possible that some kind of group behavior that was not being tested for has emerged and was not detected through any of the visual inspections or operationalisation values. Furthermore, since no test runs were performed on unfinished simulations, it is always possible that stigmergic behavior has evolved from a kind of intermediate group behavior which disappeared when the agents started to exhibit stigmergic behavior. Only a future research which does perform intermediate tests could confirm or refute this possibility. One way to do this is to run both the baseline and the experimental tests, as well as the t-tests, every time a fixed number of generations has evolved. This way, the behavior of the agents can also be analysed during the simulations, which could uncover any intermediate behaviors that the agents exhibit during certain phases in their evolution.

4.3.6 Negative t-values

Another unexpected observation in Table 3.1 is the abundance of negative t-values for the operationalisation values. A negative t-value for an operationalisation type in a setting means the corresponding operationalisation values in that setting are lower in the experimental condition than in the baseline condition in the test simulations. In other words, the amount of cooperation actually decreases if the agents are able to perceive each other. Again, no significant results were found, so this may simply be noise, but the fact that the direction of the potential effect is opposite from the expected direction is enough to warrant a discussion. In the experiment, this occurred especially often (in 7 out of 12 settings) for the clustering behavior. This could be explained by once again taking into account that agents who are near other agents may have a harder time gathering food, since other agents may take the only available food in the vicinity. It is therefore possible that agents will actually utilise being able to perceive each other to move away from each other. In Section 2.4.3, I predicted clustering behavior would be beneficial to the agents if the predators and group protection parameters were active, and the food clustering was set to a high value. In Table 3.1 this does not appear to be the case. The most obvious explanation for the agents not forming groups in these settings is that they get stuck in local maxima, as explained in Section 4.2. If no effective group behavior gets developed, it is sensible that the agents will move away from each other if they can detect each other. This would mean decreasing the clustering operationalisation value in the experimental condition compared to the baseline condition.

4.3.7 Standard Deviations

One aspect of the data that immediately springs to the eye when viewing Figures 3.2, 3.3 and 3.4, is the height of the standard deviations. Especially for the Following behavior the standard deviations are unusually high. The individual test runs are treated as single observations in the experiment, so a high standard deviation means there are large differences between the values in the different test runs. This means the results are inconsistent, and it may have been advisable to generate more different simulations and test runs. This suggests significant effects could be masked by a high standard deviation, and could potentially be uncovered by running a larger number of simulations and test runs

4.3.8 Visual Inspection of the simulations

Aside from analysing the generated data, the visualisations of a select few simulations were inspected. A number of trends were observed:

1. Agents will almost always move to and collect food that is within their sensory range. This was expected, collecting food that is within the sensory range of an agent is the most basic successful behavior the agents can develop, and increases the

fitness of the agents by a large amount when exhibited.

2. Many agents neglect to move if they can not detect food sources. These agents only exhibit the most basic of foraging behavior, they move towards food sources if they can detect any, and do not undertake any actions if they can not. This often results in a large cluster of agents remaining in the nest, since all agents who find food move to there, and any food sources in the area around the nests are quickly depleted.

3. Agents that do move out of the nest often collectively move in a single direction, with individual agents leaving the group to gather any food that they may detect under way. Occasionally, this preference for a single direction is so powerful that the agents will continue to move into that direction when they reach the edge of the field, which results in them getting stuck there. This is despite the fact that the agents have a dedicated sensor for detecting the walls at the edge of the field, which should enable them to avoid getting stuck against them.

Trend 1 is almost trivial, and will not be discussed further. Trend 2 is an indication that more advanced search mechanisms have not been developed by these agents.If more advanced mechanisms to search for food such as stigmergic trail following were developed, the agents should be able to find at least some food reliably, therefore making remaining in the nest suboptimal behavior. This would likely mean that agents who actively search for food get higher fitness values than agents who stay in the nest, which means the behavior to stay in the nest should quickly disappear over the generations. The fact that this behavior is still present during the test simulations means that effective, advanced search behavior is unlikely to have developed. Trend 3 is most interesting in the context of the researched behaviors, because the agents move collectively. However, this appears to happen in the baseline condition as well, which means the agents can independently move in groups. Most likely, the agents have developed a slight preference for a single direction through environment factors in the specific simulation environments they evolved in and genetic drift, which they then follow independently from each other. It is still possible this effect is strengthened by agents also moving towards each other or following stigmergic trails, but no significant effects for that where found. If this behavior originates from random factors or genetic drift, this is an indication that the number of agents is too small. A larger gene pool would lower the probability of a large amount of agents converging to a behavior that is not beneficial in the long run. A larger number of simulation runs might also be called for, since random environmental factors play a smaller role if more simulations were run. The fact that agents occasionally walk against the edge of the field may be caused by this not occurring often enough for the agents to evolve an avoidance behavior for the walls.

4.4 Conclusions

In the implemented environment, neither clustering, following or stigmergic following behavior has been shown to emerge in any of the investigated parameter settings. This could be due to the agents getting stuck in local maxima, an impractically large neural networks, too much abstraction in the environment, too many different experiments were run with too few test runs, or simply because the investigated parameters are insufficient for the emergence of the researched behaviors. It has been shown that in a number of settings, the fitness values decrease if the agents can detect each other and stigmergic traces, which suggests the size of the Neural Network may be detrimental to its ability to learn. A number of observations were made of the data, which may or may not be generalisable to further experiments. A lower average fitness in simulations with a high value for food clustering suggests agents do not form large groups which can take advantage of food clusters. A positive connection was observed between the amount of following performed in the experiments and the size of the food clusters. Lower, though still non-significant, p-values for the tests for stigmergic behavior than for the tests for clustering and following behavior suggest that it may be possible for stigmergic following behavior to evolve from individuals without an intermediate group phase. The observed negative t-values for clustering, following and stigmergic following suggest the agents may have learned to stay away from each other in some settings, resulting in them spreading out over the environment. The high standard deviations of the operationalisation and fitness values suggest significant effects could be masked and a more extensive research is advisable. Visual inspection of the data has shown agents to often collectively move in a single direction, although they sometimes do this as well if they are unable to perceive each other.

Although no significant effects were found in this exploratory research, the observed trends in the data can be interpreted to suggest more focused future research. Firstly, replicating this experiment but only researching a small number of parameter settings could be valuable future research. This research would be especially valuable if a much larger number of simulations per setting and tests per simulation were run. A number of relatively low p-values were found (See Table 3.1 and 3.2) but none were significant. A more thorough and focused research could unveil whether there are no effects or whether this research was too small and too wide, as suggested in Section 4.2. The Parameter that I would suggest such a future research to focus on is the Food Clustering Parameter, since this appears to be correlated to both the Fitness (4.3.3) and the performed Following (4.3.4). It would be especially interesting if future research focused on the question whether stigmergic following behavior, or swarm behavior in general, is likely to evolve directly from individuals or through an intermediate group phase. Observation 4.3.5 suggests that it is likely in this simulation that agents can develop stigmergic behavior without first going through a group phase. If a more extensive future research could confirm the presence of stigmergic following behavior in the simulations while finding no significant effect for clustering or following, it will be made likely that swarm behavior can evolve directly from individuals, without requiring an intermediate group behavior. To properly test for intermediate group behaviors, this future research would also need to run tests on unfinished simulations.

Another idea for future research is to replicate this experiment but limit the size of the neural network as suggested in Section 4.2. This research could uncover whether the size of the neural network was indeed detrimental to the performance of the agents and the ability of the network to learn more advanced behavior. If this is found to be the case, the research could also find out if the investigated group behaviors would emerge under certain parameter settings if the size of the neural network was reduced. I would suggest such a research to run two-tailed t-tests for the difference in fitness between the baseline and the experimental conditions, since the extended research in Section 3.3 showed the fitness can be reduced if agents can perceive each other.

Finally, investigating the same parameters and group behaviors as was done in this research, but in a new environment and with new agents built from scratch, could be valuable future research. As noted in Section 4.2, the choices for a task, environment and agents inevitably lead to a specific niche, which may have a profound impact on the behaviors that evolve. Many choices have to be made when building the actual implementation, which are not necessarily all part of the experimental design. A small feature in the software environment or the behavior of the agents that may seem trivial to the outcome could have a large impact. Building a new implementation from scratch and comparing the outcomes with the outcomes in this research could provide valuable insights as to what specific features have an impact on the evolution of clustering, following and swarm behavior.

Bibliography

Mark A Bedau and Paul Ed Humphreys. *Emergence: Contemporary readings in philosophy and science*. MIT press, 2008.

Robert Boyd and Peter J Richerson. The evolution of indirect reciprocity. *Social Networks*, 11(3):213–236, 1989.

Marco Dorigo, Mauro Birattari, and Thomas Stutzle. Ant colony optimization. *Computational Intelligence Magazine, IEEE*, 1(4):28–39, 2006.

Simon Garnier, Jacques Gautrais, and Guy Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.

Simon Goss, Serge Aron, Jean-Louis Deneubourg, and Jacques Marie Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76(12):579–581, 1989.

DM Hamby. A review of techniques for parameter sensitivity analysis of environmental models. *Environmental monitoring and assessment*, 32(2):135–154, 1994.

William D Hamilton. The genetical evolution of social behaviour. i. *Journal of theoretical biology*, 7(1):1–16, 1964.

Jeff Heaton. *Introduction to neural networks with Java*. Heaton Research, Inc., 2008.

Anil K Jain, M Narasimha Murty, and Patrick J Flynn. Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323, 1999.

Benjamin Kerr, Peter Godfrey-Smith, and Marcus W Feldman. What is altruism? *Trends in ecology & evolution*, 19(3):135–140, 2004.

Leslie Marsh and Christian Onof. Stigmergic epistemology, stigmergic cognition. *Cognitive Systems Research*, 9(1):136–149, 2008.

National Research Council. *The Role of Theory in Advancing 21st-Century Biology: Catalyzing Transformative Research*. The National Academies Press, Washington, DC, 2008. ISBN 978-0-309-11249-9.

Stefano Nolfi and Dario Floreano. Evolutionary robotics. the biology, intelligence, and technology of self-organizing machines. Technical report, MIT press, 2001.

Martin A Nowak. Five rules for the evolution of cooperation. *science*, 314(5805):1560–1563, 2006.

Martin A Nowak and Robert M May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.

Adin Ross-Gillespie, Andy Gardner, Stuart A West, and Ashleigh S Griffin. Frequency dependence and cooperation: theory and a test with bacteria. *The American Naturalist*, 170(3):331–342, 2007.

Erol Şahin. Swarm robotics: From sources of inspiration to domains of application. In *Swarm robotics*, pages 10–20. Springer, 2005.

Glory H Shah, CK Bhensdadia, and Amit P Ganatra. An empirical evaluation of density-based clustering techniques. *International Journal of Soft Computing and Engineering (IJSCE) ISSN*, pages 2231–2307, 2012.

Pang-Ning Tan, Michael Steinbach, Vipin Kumar, et al. *Introduction to data mining*, volume 1. Pearson Addison Wesley Boston, 2006.

Arne Traulsen and Martin A Nowak. Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952–10955, 2006.

Robert L Trivers. The evolution of reciprocal altruism. *Quarterly review of biology*, pages 35–57, 1971.

Yudong Zhang, Praveen Agarwal, Vishal Bhatnagar, Saeed Balochian, and Jie Yan. Swarm intelligence and its applications. *The Scientific World Journal*, vol. 2013, 2013.