Helping One Another: How Cooperation and Diversity Evolve

Bachelor Thesis by Saskia Koldijk (0610038)
saskiakoldijk@student.ru.nl
Radboud University Nijmegen

Supervisors: Ida Sprinkhuizen-Kuyper and Iris van Rooij

July 8, 2009

Abstract

Many animals and humans live in diverse populations in which specialized individuals help each other. From an evolutionary point of view this is puzzling: why should individuals help others if this lowers the chance that they themselves are fittest and survive? Evolutionary biologists assume that cooperation mechanisms are at work. This study investigates the hypothesis that cooperation mechanisms can cause the evolution of a cooperative and diverse population. This is done by means of simulations with an evolutionary algorithm. The results indicate that cooperation mechanisms can indeed, under certain circumstances, serve as an explanation for the evolution of cooperating and specialized individuals. The implications of these results for three different domains - Biology, Society and Robotics - are discussed.

Contents

1 Introduction

Everywhere around us we can observe that people or animals help each other and solve tasks in cooperation. There is great diversity in a population with individuals specialised on different skills. Often the individuals involved in a cooperative act can combine their strengths and can accomplish tasks more efficiently in this way.

From an evolutionary perspective however there is competition between all individuals. Only the fittest survive. Someone who helps others would lower the chance that he himself is fittest. So why would an organism help somebody? Why should he specialize on one skill if that means he is dependent on the help of others? This should theoretically lead to a homogenous and uncooperative society. The question which arises then is: How could a diverse population have evolved in which individuals help each other?

Evolutionary biologists like for example Sachs et Al. [16] describe that cooperation can evolve because there are certain cooperation mechanisms at work. Such mechanisms make sure that an individual who has costs by helping another individual gets something back for it, in the long run. A mechanism could be for example: ‘If I help you now, you help me later’. Another mechanism is that the helper directly gets part of the task reward. Can these mechanisms indeed explain the evolution of cooperation? Moreover, does the presence of cooperation mechanisms also lead to the evolution of a heterogeneous population?

Nowak [14] states the idea behind helping is that it is less costly for the donor than it is beneficial for the recipient. This is especially true when a specialist helps in solving part of a task where the individual itself is not skilled. Moreover the biologists Anderson and Franks [2] explain that solving tasks in teamwork ‘may well favor individual differences and constancy in tackling those subtasks for which the individuals are particularly well-suited.’ (p. 24). This indicates that the possibility to cooperate may cause that specialization evolves.

Certainly the task type also has some influence on the diversity of a group. The utility of diversity depends on the task [4, p. 1]. If all tasks are simple they can be solved by individuals. If tasks get more complex it can be more efficient if different individuals specialize on certain skills and then work together to solve the task.

The research question I address in this paper is: To what extent can cooperation mechanisms (in combination with certain task types) explain the evolution of cooperative and diverse populations? To investigate this I used an evolutionary algorithm. In this way I simulated evolutionary processes and analysed the development of the population. The results of this study give more insights in the underlying mechanisms of cooperation and specialization and furthermore can guide the design of cooperative robots.

In the remainder of this section I will explain some theoretical background (Section 1.1), the cooperation mechanism I am going to use (Section 1.2) and what exactly an evolutionary algorithm does (Section 1.3). The structure of the rest of the article is as follows: In Section 2 I will define the experimental method in terms of the environment, the settings of the evolutionary algorithm and the fitness evaluation. Furthermore I will outline the experimental conditions for the various simulations. In Section 3 the results of the simulations will be presented and in Section 4 I will summarize the most important findings and will give some interpretations. In Section 5 further implications of the acquired results will be discussed. Ideas for further research will be given in Section 6. The paper will end with the conclusions in Section 7.

1.1 Theoretical background

Bowles and Gintis [6] analysed the evolution of a population when a cooperation mechanism called ‘strong reciprocity’ (also referred to as ‘altruistic punishment’) is at work. The result of their study is that indeed high levels of cooperation could be sustained in the population when this mechanism is at work, which means that all individuals worked to contribute to the global goods.

In my research I do not only analyse if cooperation mechanisms can lead to a cooperative population, but I go one step further by also investigating if cooperation mechanisms can explain a diverse population. This is a connection not made before. As method I used an evolutionary algorithm instead of agent based simulations. I started with a random initial population and investigated if discriminable groups of certain specialization evolved. My cooperation mechanism was two-way byproduct benefits, which I will explain next.

1.2 Cooperation mechanism

Sachs et Al. [16] give a framework for cooperation mechanisms. One of the cooperation mechanisms is called ‘two way byproduct benefits’, which includes synergism: ‘actions or coordinated behaviours that are automatically more profitable when performed in groups’ [16, p. 145]. In this setting everyone benefits from working together. The idea is that each individual can solve the part of the task where he or she is skilled in. The helping individual gets part of the reward. The advantage for the other individual is that she does not have to do subtasks that she is not good at and thus would be very costly. Indirectly both individuals benefit each other (see Figure 1, right).

In my experiments I will compare these settings with cooperation mechanisms with the setting where no cooperation mechanisms are at work. In that case the first individual gets the reward of the whole task. The helping individual does not get any reward, which means that he cannot improve his fitness. He will end up with costs (see Figure 1, left).

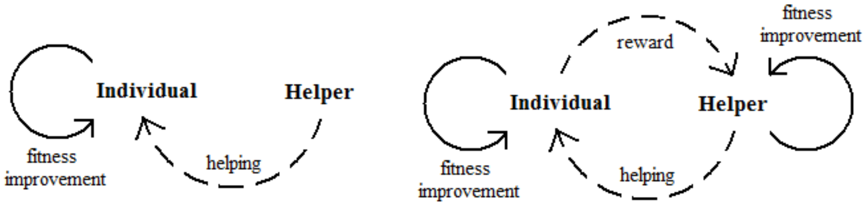

Figure 1: Left: No cooperation mechanism, right: Cooperation mechanism: two-way byproduct benefits

1.3 Evolutionary Algorithms

An evolutionary algorithm is a computational program inspired by natural evolution [7]. The basic idea behind it is that the program should find a solution that best fits some criteria. It does this by simulating evolutionary processes. The basis is the evolutionary cycle, consisting of the following steps: initialization, parent selection, recombination, mutation, survivor selection and finally termination (see Figure 2).

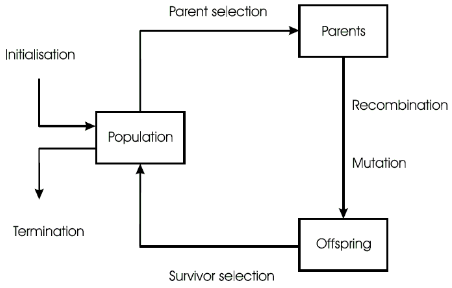

Figure 2: General scheme of an Evolutionary Algorithm (after Eiben and Smith [7])

The computer program starts with some random solutions, called individuals, which form the initial population. How ‘good’ an individual is can be determined according to the so called fitness function. This is a function which assigns a value to an individual, which indicates how good he fits the requirements. In each generation only the best individuals will create offspring. The genes from two parents can be recombined to create new individuals. There also can be mutation of certain genes, which means that some randomness can be added to the genes of new individuals. In the long run only the best genes will survive. In this way a population can be evolved which suites the given criteria (fitness function) best.

Several articles report the success of evolutionary algorithms for evolving cooperative behaviour. Experiments by Suzuki and Arita [17] have shown that cooperative behaviour can evolve and that it is a necessary requirement for teams to perform well, as they illustrate for a simple form of robot soccer. Furthermore Quinn et al. [15] report that artificial evolution was able to evolve a team in which robots maintain functionally distinct roles in order to achieve a task.

One of the fundamental factors for evolving a cooperative team is the credit assignment or ‘how to fairly split the fitness of a team to all of its participants’ [9, p. 23]. Should a helping individual also get part of the task reward? And if yes, how much? As a study about reinforcement learning indicates: When the entire team is jointly rewarded, teams tend towards heterogeneous behaviour [3, p. 7]. In this way the group achievement is considered instead of single individual performance and the team members are encouraged to adapt to each other. Jointly rewarding the individual and her helper could thus mean that a diverse population evolves: the individuals might specialize on different skills and trust on the help of other specialized individuals to solve the tasks. The particular fitness evaluation used (which incorporates the cooperation mechanism) is thus the most important part of my evolutionary algorithm.

2 Method

In this paragraph I will describe the fundamental aspects of the experimental method: The environment, the evolutionary algorithm, the fitness evaluation and the experimental conditions.

First I will explain the environment which is defined in terms of the task, the population and the encounter (Section 2.1). Next the exact evolutionary algorithm settings will be given (Section 2.2). The fitness evaluation will be explained in terms of rewards and costs associated with the tasks and skill levels (Section 2.3). Finally I will explain the experimental setup in terms of the different conditions (Section 2.4).

2.1 The environment

2.1.1 The task

Inspired by the article of Bowles and Gintis [6] I have chosen to keep the task abstract. That means that there is not one specific task, but a task is only defined by the skills which it requires to do it. For example a task could require skill A and B, but not skill C, represented as (1 1 0). A task can be accomplished alone or, probably more efficiently, by working together. The different skills can be fulfilled by different individuals, thus the definition of a team task by Anderson and Nigel [1] applies: ‘A team task requires different subtasks to be performed concurrently for successful completion. There is a division of labor within a team.’ (p. 534)

By keeping the task abstract, the experiment and the results are not specific to one situation but apply to all tasks that can be formulated in such a way. As illustration I will give an example: Imagine the task is to drive by car and read the map to find your way to an unknown destination. The skills are map reading and car driving. You might be good in map reading, but have no drivers’ license. You can ask someone to help you to fulfil the task by driving. Maybe you have a drivers’ license and also have the skill of map reading. But nevertheless it will probably be better and more efficient to fulfil the task with a partner, so that one person can focus on driving and the other on map reading.
I use the following task types in my simulations:

- There are be simple tasks that can be fulfilled by individuals alone:
  - tasks where one skill is required (‘Simple tasks’): (A), (B), (C)
- Other tasks are more difficult and might be solved more efficiently when working together:
  - tasks where different skills are required (‘Different tasks’): (A B), (B C), (C A)
  - tasks where the same skill is required more than once (‘Same tasks’): (A A), (B B), (C C)

2.1.2 The population

As the tasks require different skills, the individuals in the population are defined by the degree to which they posses these skills. There are 3 skills: A, B and C. Each skill gets associated with a number between 0.1 and 1, where 0.1 means being bad in the skill and 1 means being perfect. It is assumed that no individual can totally lack a skill (skill level = 0). Being skilled means that doing this aspect of the task has a low cost. To prohibit that all individuals just evolve towards ‘perfect in everything’ a maximum of 1 is imposed on the sum of the skills. In nature there are also often such constraining trade-offs, for example an individual specialized on prey A loses the ability to efficiently capture prey B [5, p. 14]. A maximum of 1 means that an improvement of one skill always is at the costs of other skills. In the experiments I assume a total skill potential which is identical for all individuals. It could be argued that some people have a higher skill potential than others. But as this research is about diversity it is more important how the skills are present in relation in each individual than the total skill amount of individuals. So for simplicity reasons the total skill potential is fixed at 1.

2.1.3 Encounters

For assessing the fitness of the individuals in a population there are encounters, which are described next. There are tasks that need to be fulfilled. An individual is chosen and randomly assigned a task. The individual may choose to fulfil the task alone or she may ask for help. When she asks for help N individuals are randomly selected from which she can choose one as a partner. If none of the individuals is suited she can choose to fulfil the task alone. When she chooses a partner they fulfil the task in team work. In this way the teams are spontaneously formed as tasks demand it. This is in line with what Anderson and Franks say [2, p. 26].

After each individual in the population solved M tasks alone or with help the encounter is finished and the fitness is determined. The exact manner of fitness assessment depends on the cooperation mechanism, but generally a task has an associated reward and the individuals will have costs in performing the tasks, depending on the degree to which they posses the skills. The fitness will be the reward minus the costs summed over the M tasks. Furthermore it is possible to gain extra fitness. This can happen when an individual acts as helper. More details about the fitness assignment follow in Section 2.3.

Some decisions had to be made about the exact course of the encounters. These details follow in the next paragraphs.

When is cooperation possible? When a task has two subtasks it can be divided between two individuals and solved in cooperation. When the task only has one subtask it cannot be divided, so cooperation is not possible.

How is the task divided? If a task will be solved in teamwork the question arises how the task should be divided. I will illustrate this with an example:

Assume the task requires skills B and C. The individual herself has the skill pattern (0.1, 0.4, 0.5) and the potential helper has the skill pattern (0.1, 0.1, 0.8). There are different divisions possible: the first individual could solve task C in which she is most skilled, but this would mean that the helper must do skill B in which he is very bad. Thus this would not be a very effective choice.

In my experimental design I chose to divide the task as effective as possible, i.e. the task is divided such that the overall costs are minimal. In the example the most effective division is that the first individual does B and the helper does C. It should be noted that this could mean that an individual might have to solve a subtask for which his skill level is not highest.

How many individuals are chosen as potential helpers (N)? In most real world environments you will have a limited number of people that are around to help you. In light of this, choosing from the whole population as potential helpers would not be realistic. Moreover if the individuals are free to choose anyone to help, this could lead to cases where always the same individual is chosen as helper, for example if the skill A is rare. As a consequence the individual skilled in A would have the chance to gain much extra fitness as doing skill A is easy and cheap for him. A high fitness means that he will have a lot of offspring. This could cause the whole population to flip, which means that there would be too many individuals with skill A and another skill might lack which would cause further flips. No stable and analysable state could be reached to answer my research question.

Directly assigning one individual as potential helper does also not seem a good choice, because in that case it is quite likely that a proposed helper cannot help you. Most ideally there must be some possibility to choose a helper who is skilled in things you are not. It seemed reasonable to me to select 3 individuals randomly as potential helpers. In the real world these might be the persons around you. But there is no reason not to choose 4 or 6 individuals. This is a choice I made and I will keep this variable fixed at 3 for reasons of simplicity. Small variations in the number of chosen individuals showed to have no influence.

Which individual will she choose for helping her? The individual will choose the helper with whom she will have the biggest fitness gain. She takes in mind that a helper will not help if his fitness gain is negative. So the individual will check the following before she makes her decision:
- What is my fitness gain if I solve the task alone?
- If I work together, given the best task division what is my fitness gain and what is the fitness gain of my helper?
- If the fitness gain of my helper is negative then he will not help, so there is no need to ask him.
- From the remaining helpers she will choose the one with whom her own fitness gain is biggest; she might as well decide to work alone if that yields the best fitness gain for her.

When does a chosen individual help? An individual is only asked when his fitness gain will be positive (helper threshold = 0), so a chosen individual always helps.

How many tasks must an individual solve (M)? If only one task would be used for fitness assessment, the fitness would be very dependent on the randomly assigned task. An individual could have luck and be perfectly suited for the assigned task. This would lead to high fitness, although the individual might be bad in solving other tasks. To prevent this, each individual has to solve 10 tasks. In this way the general potential of each individual is measured. Using more tasks would make the assessment even better, but also computationally more intensive.

Must a task be accomplished? As most tasks in real life are obligatory, it seems most realistic that the individuals must accomlplish their assigned task. Moreover the possibility to not accomplish a task at all would unnecessarily complicate the experimental setting. For these reasons all tasks must be solved.

2.2 The Evolutionary Algorithm

In the following I describe the settings I chose for each of the stages in the evolutionary cycle (see Figure 2). Table 1 provides an overview. All these settings are quite standard evolutionary settings [7]. I used ECJ 18, a Java-based evolutionary computation research system, to implement the evolutionary algorithm [12].

- The Representation: The individuals are represented as real-valued vectors. There are three genes with values between 0.1 and 1 which are normalized such that their sum equals 1. These indicate the skill levels for the three subtasks.

- Initialization: A population of 100 individuals is used. All individuals are initialized randomly (a uniform distribution between 0.1 and 1 is used). A normalization procedure ensures that the sum of the three numbers is equal to 1. Normalization might in some cases lead to values a bit smaller than 0.1, but this has no important consequences for the algorithm.

- Parent selection: The parent selection is tournament based with a tournament size of 3. That means that 3 individuals are randomly selected and the individual with the highest fitness value will be chosen for reproduction. With this method 100 parents are selected, who in total produce 100 individuals offspring.

- Recombination: As recombination method one-point crossover can be used, but as the representations are quite short, I chose not to use recombination, but only mutation.

- Mutation: With a chance of 1/3 (= 1/genome length) mutation is applied to a gene. As mutation operator non-uniform mutation is used. This means that an amount drawn randomly from a Gaussian distribution (N(μ, σ) with $\mu = 0, \sigma = 0.05$) is added to a gene. It is ensured that no gene can be lower than 0.1 and higher than 1. A normalization procedure will ensure that the sum of the three numbers stays equal to 1.

- Survivor selection: There are different methods for survivor selection. For example one could choose to let only the best individuals (parents or children) survive. In these experiments the survivor selection is generational, which means that all the generated 100 children survive.

- Termination: The individuals in the population will adapt more and more to the requirements defined in the fitness function. In these experiments the evolutionary algorithm terminates after 200 generations. In pilot studies I observed that there is no great fitness improvement after the first 10 to 20 generations, so by running 200 generations all important (long term) effects could be observed.

Population size	100
Length of genome	3
Minimum value of genes	0.1
Maximum value of genes	1
Parent selection method Tournament selection with size	3
Recombination	-
Mutation probability	1/3
Mutation type	Gaussian, $\mu = 0, \sigma = 0.05$
Termination	After 200 generations

Table 1: Evolutionary settings

2.3 The Fitness Evaluation

In contrast to the quite standard evolutionary algorithm settings explained above, the fitness evaluation is specifically written for this experimental setting. As already mentioned the fitness of the individuals is determined in an encounter, where each individual is randomly assigned 10 tasks, which it has to solve. Per task some fitness can be gained. Moreover an individual might gain extra fitness if he helps.

If a task is solved alone the contribution to the total fitness of the individual will be the reward of the task minus the costs:

$$Fitnesscontribution_{ind} = Reward(task) - Cost(task, skills) \tag{1}$$

When two individuals work together the task is split up into two subtasks. The fitness contribution for each individual depends on the proportion reward that the helper gets (pr). These formulas are comparable to those used in [11].

$$Fitnesscontribution_{ind} = ((1 - pr) \times Reward(task)) - Cost(subtask_{ind}, skills_{ind}) \tag{2}$$

$$Fitnesscontribution_{helper} = (pr \times Reward(task)) - Cost(subtask_{helper}, skills_{helper}) \tag{3}$$

I had to make a choice how much of the reward the helper exactly gets, because this was not stated in the model. I chose to use two settings: In the first setting the helper gets 25% of the task reward (pr = 0.25), which means that help is quite cheap. In the second setting the helper gets 50% (pr = 0.5), which means that help is expensive.

Reward The more skills a task requires the bigger is the reward. This can be expressed by the function:

$$Reward(task) = 10 \times Sum(subtasks) \tag{4}$$

For a difficult tasks which requires skills A and C this yields: $Reward((1, 0, 1)) = 10 \times sum(1, 0, 1) = 20$

Cost The more skilled an individual is the less does the subtask cost him. There are different choices to express this as a function. The function could for example be linear, but I chose to use a concave function. I made this choice, because this function reflects the properties of a skill best: If you are very bad in a skill it will cost you extreme much to use it. If you get better in the skill this will soon save you many costs, but from being good to being perfect in a skill only a few costs will be saved. Note that Bowles and Gintis [6] also used a concave cost function in their experiments. As function this can be expressed by (see Figure 3):

$$c(skilllevel) = \frac{1}{skilllevel} \tag{5}$$

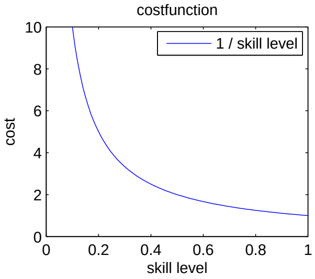

Figure 3: Costfunction

The total cost for a task is determined according to the following function:

$$Cost(task, skills) = \frac{subtaskA}{skilllevelA} + \frac{subtaskB}{skilllevelB} + \frac{subtaskC}{skilllevelC} \quad (6)$$

If an individual chooses to solve a task alone there are extra costs. This is because a task can be solved less efficiently alone. This is implemented by multiplying the cost by a factor. In a pilot study several values between 1.1 and 1.7 were tested (for more details, see Appendix Section A). To avoid making the experiments even more complex I decided to make this parameter fixed. Choosing an intermediate value seemed the best and also most realistic option, so this parameter has the value 1.3.

To illustrate, for an individual with the skill levels (0.2, 0.3, 0.5) who solves the task (1, 0, 1) this yields: $Cost((1,0,1),(0.2,0.3,0.5)) \times 1.3 = (\frac{1}{0.2} + \frac{0}{0.3} + \frac{1}{0.5}) \times 1.3 = 7 \times 1.3 = 9.1$

2.4 Conditions

In the experiment two variables are manipulated: The type of the tasks and the cooperation mechanism, which is implemented as what proportion of the reward the helper gets (pr). Each variable has 3 possible values, which yields a 3 by 3 design. The levels of task type are: Simple (one skill required), Different (two different skills required) and Same (two times the same skill required). The levels of cooperation mechanism are: none (pr=0%), two way by-product benefits (pr=25% and pr=50%).

3 Results

In this section I describe and explain the most important results of the simulations. First of all, the hypothesis that cooperation mechanisms actually lead to cooperation between the individuals requires verification. Cooperation mechanisms make cooperation possible, but that does not automatically imply that individuals actually do cooperate. After that I will turn to the question how the population is spread at the end of the evolutionary process. I will investigate if the population is homogenous or heterogeneous and specialized. Furthermore I will analyse the average group fitness to find out in how far solving all tasks in cooperation in a diverse group is more efficient than solving all tasks alone in a heterogeneous group. Finally I will analyse how stable the acquired results are across generations.

The section is subdivided into the different measures used: Amount of cooperation (Section 3.1), heterogeneity and specialization (Section 3.2), stability of the results (Section 3.3) and average fitness (Section 3.4).

3.1 Amount of cooperation

To investigate if cooperation mechanisms actually lead to cooperation the number of cases in which individuals solve tasks together is stored (see Figure 4). Furthermore this measure could be used to investigate if working together correlates with heterogeneity of the population.

No cooperation mechanism In the case where the helper does not get any reward there is no working together at all. This is due to the fact that a potential helper would only make costs in this setting. As individuals only help if they benefit, no one will help, so no cooperation evolves.

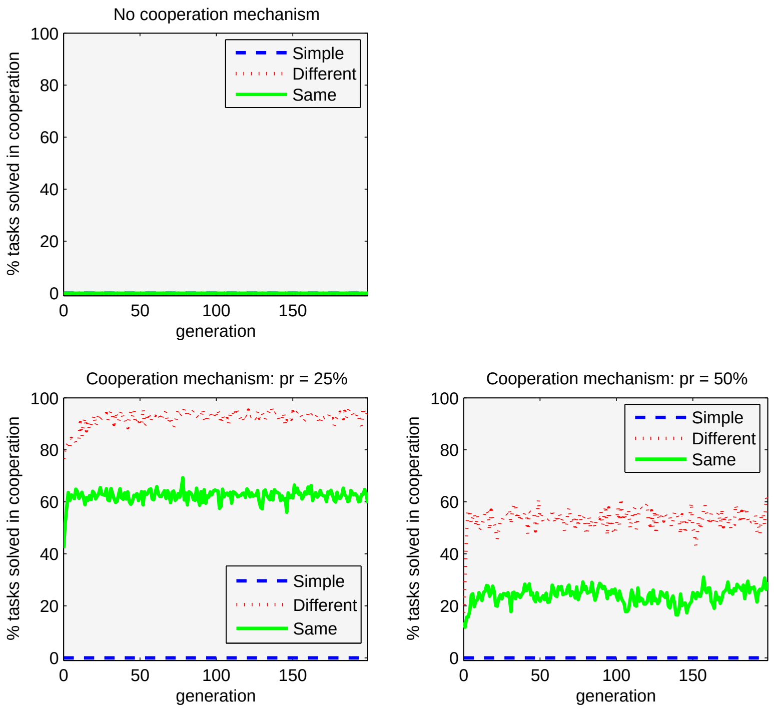

Figure 4: Amount of cooperation

Two way byproduct benefits When the helper gets 25% of the task reward, helping does take place in certain cases: For the simple tasks there is no working together. This is due to the fact that a simple task only has one subtask, which means that it is not possible to work together. For the Different tasks the first population already works together in 80% of the cases. Over the first 20 generations this still improves until a quite stable amount of above 90% working together is reached. This value is the highest reached over all simulations. The amount of working together for same tasks is also highest in this setting. In the first generation 40% of all individuals work together, which quickly improves to above 60% in the first few generations. After that the value stays quite stable.

In the next setting the helper gets half of the reward. Again for simple tasks there is no working together as this is not possible. For different tasks the amount of working together quickly reaches about 50 to 60%. On same tasks only 20 to 30% of all individuals work together.

3.2 Heterogeneity and Specialization

The main dependent variables in the experiments are the degree of heterogeneity and the amount of specialization of the group. This can be visualized by plotting the skill vectors of each individual. The more they are spread, the more heterogeneous the group is. The specialization can be measured by the amount of clustering, i.e. the degree to which homogenous subgroups are formed in the population. Interesting is also the change in heterogeneity and specialization as the group evolves, visualized as a movie (available on request).

Figure 5 shows the begin population. On the x-axis skill A is plotted and on the y-axis skill B is plotted. As all three skills sum up to 1 the resulting skill level of skill C can also be read from this plot: points at the left bottom of the plot represent individuals with low skill levels on A and B and thus a high skill level on skill C. All individuals are initialized randomly which yields a very heterogeneous, but not specialized start population. You can see that no individual can lie outside of the triangle, as the sum of all three skills must be 1.

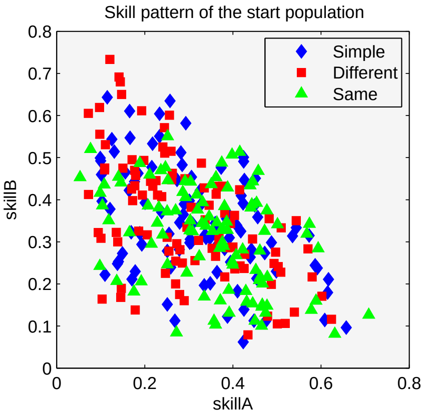

Figure 5: Skill pattern of the first generation

In the plots of Figure 6 the population of the last generation is pictured.

No cooperation mechanism In the first situation of Figure 6 the helper did not get any reward. All individuals converged to the middle. This means that all individuals are equally good in all three skills. The population is very homogenous.

We already saw in Figure 4 that in this case there was no working together at all, so this might be the first indication that no cooperation leads to a homogenous population without specialization. As an individual cannot count on help he probably tries to be equally good in everything to best prepared to solve any possible task.

Two way byproduct benefits The second plot shows the case where the helper gets 25% of the reward. Here an interesting pattern of spreading of the population is visible. In the case of simple tasks the individuals are again quite centred around the middle, but the individuals for different tasks lie in the corners, which means that they have specialized on solving one subtask. Three groups can be distinguished which each specialize on a different skill. The individuals for same tasks lie in between and thus show a weaker form of specialization.

In Figure 4 we saw that for this setting the amount of working together was highest. Especially for different tasks 90% of them where solved in teamwork. This might mean that when there is a high possibility of solving a task together, the population is more heterogeneous and different groups of individuals specialize on different skills. For same tasks the percentage of cooperation was a bit lower, which might explain why these individuals are less specialized.

For the third plot the setting is that the helper gets half of the task reward. A quite different pattern can be seen in this case: The population for the simple tasks lies in the middle again, but the population for different and same tasks is pulled towards the axes. A point just above the middle of the x-axis means that this individual is very bad in skill B, but as good in skill A as in skill C. So the individuals do not specialize on one skill, but they seem to unlearn one skill. There is no difference between different and same tasks in this case. This is unexpected, as Figure 4 showed that the amount of cooperation is different for different and same tasks (ca. 50% and ca. 25% of the tasks are solved together respectively).

3.3 Stability of the results

As evolution is a dynamic process it is interesting to analyse how the process develops over generations. The question that arises is: Is the analysed pattern of specialization for the last generation stable across generations? And if this is the case: How long does it take until this pattern evolves? To investigate this I analysed videos of the development of the heterogeneity and specialization of the population. It turned out that the found pattern of specialization evolved after only about 10 to 20 generations. From that point on the individuals in the population do shift a bit, but the evolved groups persist.

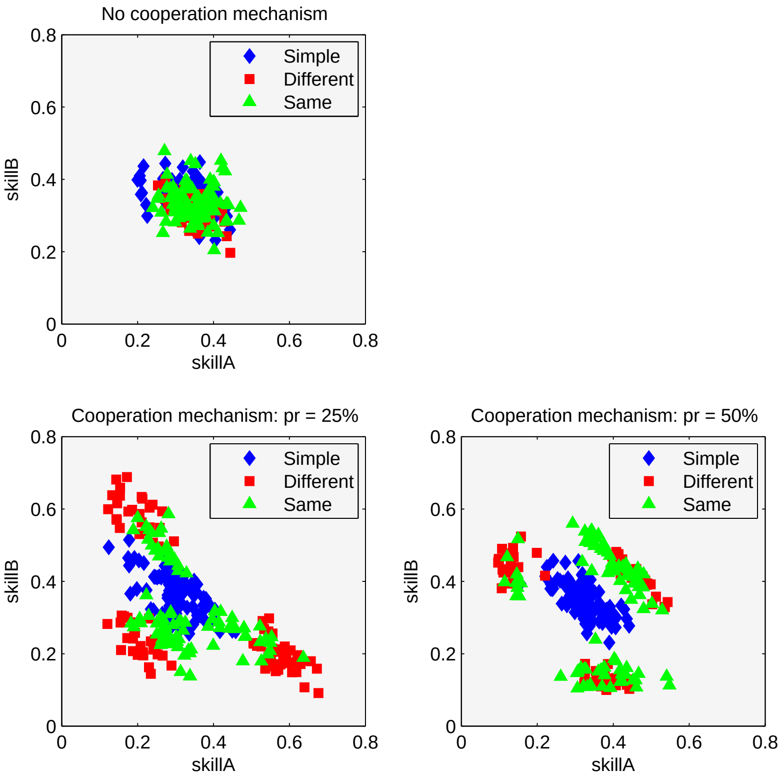

Figure 6: Skill pattern of the last generation

3.4 Average fitness

To analyse the development of the population the program keeps track of the average fitness of the population per generation (see Figure 7).

No cooperation mechanism In the first case the helper does not get anything of the task reward. In the plot you can see that already after the first few evolutions stable average fitness values are obtained. For simple tasks this value is 61 and for different and same tasks this value is twice this value: 122. The reason why the curve for simple tasks is lower, is that the associated reward for simple tasks is only half the reward as for different and same tasks (respectively 10 and 20). For computations see Appendix Section B. As we know from Figure 4 there is no working together in this case. So this condition can be seen as a baseline for the average group fitness.

Two way byproduct benefits In the second plot the results for the case where the helper gets 25% of the reward are shown. Also in this case the values settle very quickly. For simple tasks the average group fitness is again 60, which can be explained by the fact that there is no more efficient way of solving simple tasks as there cannot be any task division. Furthermore there is an improvement for as well different as same tasks in this setting: Same tasks have an average group fitness of about 145 and same tasks have a group fitness of about 135. These values are the highest compared with all other settings.

These results suggest that specializing on different subtasks is a good way to improve the performance of a population, given there is possibility to work together. A proportion reward of 25% for the helper seems an optimal setting.

In the third case the helper gets 50% of the task reward. For simple tasks the average group fitness is 61 again for reasons just explained. For same tasks the average group fitness is 122, as in the first, baseline plot. For different tasks there is an improvement compared to the baseline plot: The average fitness is 135 now. This is a quite interesting result: In Figure 6 it turned out that same and different tasks lead to the same

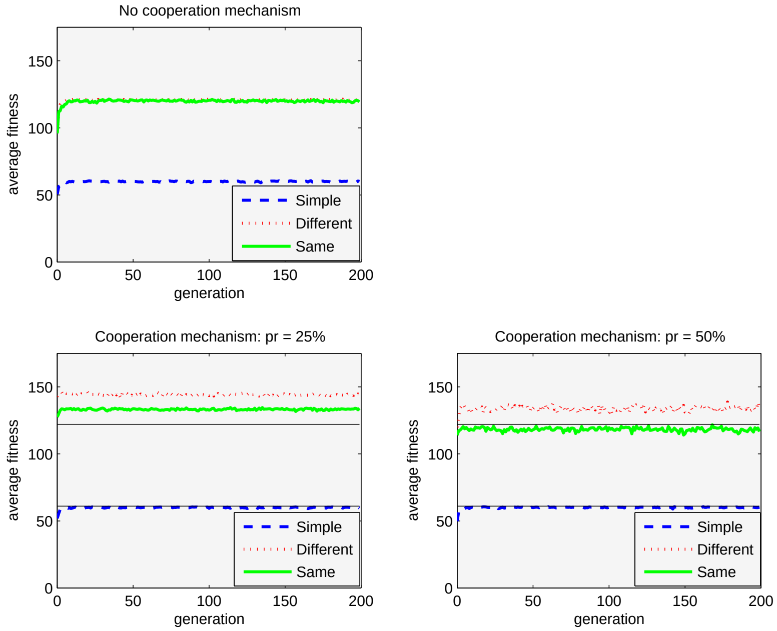

Figure 7: Average fitness

pattern of specialization. Nevertheless the efficiency only improved for different tasks. This finding might be connected to the fact that for different tasks there is more working together. It seems that the evolved pattern of specialization can only lead to more efficiency in combination with a high amount of cooperation and tasks requiring different skills.

4 Summary and interpretation of the results

In this section the acquired results are summarized and interpreted. The situations with and without cooperation mechanisms are addressed (Section 4.2 and Section 4.1). Furthermore the stability of the results is explained (Section 4.3).

4.1 No cooperation mechanism

In the setting without cooperation mechanism, the individuals did not work together at all. This indicates that without cooperation mechanism no cooperation could evolve. Moreover the population plot showed that all individuals converged to being equally good in each of the skills, independent of the task type. As expected each individual tried to be prepared for all possible tasks, as they could not count on help from others. We can conclude from this: no cooperation mechanism implies no cooperation, which in turn leads to a homogenous population.

4.2 Two way byproduct benefits

Cooperation When there is a cooperation mechanism at work the individuals actually work together in many cases to solve the dividable tasks (different and same). This indicates that cooperation mechanism can indeed cause that the population cooperates. The percentage how often individuals exactly cooperate depends on the task type and the proportion of the reward that the helper gets.

Task type - In general there is more cooperation for different tasks than for same tasks. The difference in amount of cooperation can be explained as follows:

- For different tasks different skills are required. In this case two individuals can combine their strength: each individual solves the subtask he is good in. As he can choose from 3 individuals an optimal task division can be made. This makes working together much more efficient than working alone as an individual has a restricted total skill potential, thus cannot be good in everything. For a sample calculation, see Appendix Section C.

- For same tasks two times the same skill is required. So in this setting the first individual has to fulfil the skill at least once. The potential gain by working together lies in choosing an individual which is better in this skill to support him. When the first individual is bad in the required skill there is much gain if another individual can help, but when the first individual is herself good, she can easily solve the task alone. For a sample calculation, see Appendix Section C. The chance that the individual herself is good in the required skill is relatively large (1/3), which explains why there is in total less cooperation for same tasks.

We can conclude from this that the simple presence of cooperation mechanisms does not guarantee cooperation by itself. The kind of tasks that should actually be solved plays an important role. Especially in environments were tasks require different skills the evolution of cooperation is likely.

Reward - Moreover there is more cooperation when the helper gets a low reward than when the helper gets a high reward. This can be explained as follows:

- A potential helper will be quite easy to find in both settings, because the helper does get part of the task reward. By fulfilling a subtask in which he is skilled he can easily gain some extra fitness. The decision if cooperation actually takes place lies in the hands of the first individual. She will only decide to solve the task together, if this yields her a larger fitness gain than working alone. She has to choose between doing all costly subtasks alone (plus extra costs for working alone) or giving part of the reward to a (specialized) helper.

- In the setting where 25% of the task reward goes to the helper, letting someone else do things you are not skilled in is a cheap option. So all in all as well the first individual as the helper are better of by working together. This is thus a very cooperation facilitating setting.

- The reason that there is less cooperation in the setting where the helper gets 50% of the task reward lies in the fact that too much of the reward goes to the helper. In many cases the first individual will prefer to work alone, as help is too expensive.

We can conclude from this that just the fact that the helper also gets part of the reward is not sufficient for cooperation to take place in a high level. The exact amount of the reward plays an important role.

Heterogeneity and Specialization Furthermore there seems to be a connection between the amount of cooperation and the degree of heterogeneity and specialization of the population: If there is much cooperation, the group is also more divers and specialized. How exactly the specialization of the population looks like depends on how much the helper gets. This can be explained as follows:

- Based on the fact the individual can choose from three potential helpers, there is a high probability that one of the potential helpers is specialized in a required skill. Combined with the fact that the helper gets part of the task reward there is a high chance that an individual is actually willing to help.

- Specializing on one skill, with the consequence of unlearning the other two skills is a good strategy in the setting where help is cheap: If the given task requires your specialization you can solve it easily and if an extra other skill is required there will probably be someone that helps you. As you yourself get most of the task reward this case is ideal for you. This means that more cooperation between specialists can be expected in environments where helping each other is cheap.

- In the setting where help is expensive being quite good in two skills with the consequence of unlearning one skill is the best strategy. It would be too risky to specialize on only one skill as needing a helper is quite costly. So in environments where help is expensive less specialized and cooperative individuals will be found.

Average Fitness Moreover the experiments showed that higher average fitness values can be observed in more cooperative populations. This means that helping other individuals improves the fitness of the whole group, which means that from an evolutionary perspective cooperation is favorable, as a cooperative population will outperform an uncooperative one. (This might be the reason that cooperation mechanisms evolved.)

Summary To summarize all these results: In general cooperation mechanisms can cause cooperation, which in turn leads to a heterogeneous, specialized population. The exact degree of cooperation and the specific pattern of specialization depend on the task type and if help is cheap or expensive.

4.3 Stability of the results

The experiments showed that distinguishable groups with different skill patterns were formed which persisted over generations. Normally genetic drift causes that all individuals evolve towards one optimal solution. The question is then: which is the optimal solution in this setting? Or is there more than one optimal solution? The statement by Nettle [13] may be considered here: ‘Behavioural alternatives can be considered as trade-offs, with a particular trait producing not unalloyed advantage but a mixture of costs and benefits such that the optimal value for fitness may depend on very specific local circumstances.’ (p. 625). In these experiments specializing on one skill means that the individual is worse in the other two skills, but whether this particular specialization is fitness enhancing depends on the other individuals, too, as cooperation can take place. Imagine that all individuals specialize on skill A. Then no-one can easily help with solving tasks which require other skills, which is bad for the fitness of the whole population. This indicates that the individuals must adjust their specializations to each other. As Eiben and Smith [7, p. 222] write: The fitness of an organism is entirely determined by the environmental niche in which it lives. The characteristics of this niche are predominantly determined by other organisms. Depending on the skills of the other individuals a particular own skill pattern is optimal. Hence there is coadaption between the individuals in the population and this can explain that an arrangement of three optimal skill patterns could evolve and was maintained.

5 Discussion

The results indicate that cooperation mechanisms can indeed, under certain circumstances, serve as an explanation for the evolution of cooperating and specialized individuals. The experimental setting was kept very simple, but importantly it shared essential features with solving tasks in the real world. The encounter mechanism was constructed to reflect natural, real world cooperation settings. Furthermore all tasks were kept abstract and were only defined in terms of essential properties. This set up has the advantage that the results generalize to any environment in which individuals posses different skills and tasks require some proper subset of these skills. These features hold for many real world situations. The implications of this research span several different domains, including Biology, Society and Robotics. I will next discuss these implications for each of the domains separately.

5.1 Biology

The simulated evolution in this study can be seen as capturing selection principles that may also be operational in biological evolution. The experiment showed that especially for more complex tasks, requiring the combination of several skills, much specialization and cooperation evolved under the assumption of cooperation mechanisms. The same phenomenon is also observed in many insect societies, where individuals cooperate to tackle difficult and complex tasks, such as nest building or overpowering prey. This indicates the biological plausibility of the effect of cooperation mechanisms.

According to Anderson and Franks [1], insects exploit a new organizational method. They form highly cooperative and relatively unselfish groups of individuals who are together able to tackle tasks far beyond the abilities of the individuals alone (p. 534). The simulation research can give us insights in the interrelations in biological evolution: cooperation mechanisms combined with the fact that the insects face complex tasks could be the basis of the evolution of this new organizational structure.

Furthermore Anderson and Franks [1] state: ‘Specialization, learning, and overall enhanced performance efficiency will enhance the ergonomic efficiency of a team.’ (p. 538). This means when work is done in a team, different individuals can focus on certain subtasks, which they will perform more and more efficiently due to learning. This leads to specialization of the individuals on the long run. The adaptation of the individuals observed in the simulation might thus not only be equalled with the adaptation by inheritance of favourable properties but also with the adaptation by learning processes. On the basis of cooperation mechanisms adaptation does not only take place over several generations but also within one generation.

All in all this research was able to yield a better understanding of biological phenomena.

5.2 Developments in the society

Often, the word ‘evolution’ refers to biological evolution of species under selection pressures for procreation. However, many things can evolve, among them societies of cultural habits and skills. In a more abstract sense evolutionary processes of selection are also operational in societal evolution. Instead of inheritance of properties, favourable properties prevail because people adapt their behaviour: If one sees that a particular property or skill leads to success one might also invest in this skill. In this sense the best skills or properties still survive in the population, albeit it not passing on of genetic materials but instead by adapted or learned skills.

Moreover, as in biological evolution, some coadaption has to take place in society as well: People have to adapt to their social environment, which consists of other people adapting to that same environment. How favourable some property is also depends on the properties of others. Hence, from this point of view the evolutionary processes studied in the simulation experiments can be compared with evolutionary processes in society.

Societies existing centuries ago were most often very simple and small. People lived together in a village and there was a small degree of specialization and cooperation. Over time the society circles went bigger and bigger: People band together in cities, countries and even over the whole world. The tasks encountered by individuals nowadays are very complex, i.e. requiring many subtasks and different skills. Take for example the production of a car: You need a designer, a technical expert, workers and many more. Also the range of possibilities for cooperation is large. For instance, technological developments made it possible to even cooperate with people at the other side of the world. In this way productions from China can be combined with parts from Germany and so on. This yields much more possibilities. People over the whole world can specialize on things they are good in and cooperate with each other. By this organization on a higher level much more can be accomplished and progress can be made. One of the fundamental factors is the possibility to cooperate, as the experiments indicate. Imagine you would not be able to cooperate in today’s society. On your own you would never be able to reach such a high living standard. Cooperation is thus a very strong mechanism to combine strengths to reach a higher level of living. As Nowak [14] states: cooperation makes evolution constructive and is thus the secret behind the open-endedness of the evolutionary process.

Besides the specialization on certain skills even different classes developed in the society: e.g. workers and academics. Henrich and Boyd [10] speak of stratified societies with more economic specialization. They say that stratified human societies are more like an ecosystem in which different, partially isolated cultural groups evolve mutualism as different species do. Like in the experiments, there is cooperation between groups with different properties, which need each other. Henrich and Boyd [10, p. 15] further state that more stratification will lead to higher average production than in egalitarian societies. In light of the experiments this can be explained, as a stratified society can do or reach more by the cooperation of groups with different skills.

The development towards more divers and specialized societies is also reflected in the development of an individual in his lifetime. First education is very general, but over time individuals are encouraged to specialize and advance their talents, as cooperation is possible and consequently it is not required that everybody is good in everything. This specialization implies that you must trust on the help of others. As the cooperation networks get bigger and bigger nowadays people all over the world depend on each other. As long as the situation stays stable there is no problem: Everyone needs others to survive, so cooperation is sustained as a consequence of this coadaption.

As a conclusion we can state that the results of the experiments yield a better understanding and more insights in the development of the society.

5.3 Robotics

The end state of the population in the experiments can be seen as the optimal population given certain environmental constraints. The insights gained in the experiments can be used in the design of multi-agent systems. Some important conclusions are the following. For tasks requiring only one basic skill (what we called 'simple tasks') homogenous teams perform best, but when the tasks have more subtasks and especially when they require different skills, heterogeneous and cooperating teams perform best. This relation between task complexity and heterogeneity of groups was also conjectured by Balch [4]. In settings were robot teams are required to solve complex tasks, cooperation and specialization of the robots can be facilitated by implementing a cooperation mechanism, and then evolving their control structures under selection pressures on task performance. The ability to *evolve* specialized control structures for robots is very advantageous. In this way human designers need not define the specific behaviours of the robots for each problem [17]. The way I set up the simulated evolution process can thus inform, and supplies a method for, multi agent design.

6 Further research

The experiment makes the assumption that the environment does not change in the course of evolution. The type of tasks that must be solved stays stable. If you can rely on such a stable environment it is a good strategy to specialize on some skill. Biological studies confirm this (eg. [8]). In the real world however the environment does change. In that case less specialization might be advantageous, as specialization would make you very vulnerable to sudden environmental changes. In further research the effects of an unstable, changing environment on the degree of specialization could be analysed.

Another simplifying assumption in the simulations was that all individuals only help when they can gain some fitness (helper threshold = 0). In the real world however some individuals seem more cooperative than others, so each individual might have his own helper threshold. In further research the evolution of this threshold could be investigated in combination with the cooperation mechanism 'partner choice' [16]: Individuals who are more cooperative, i.e. are even willing to make costs by helping, are preferred. This implementation might lead to an even more realistic setting. On the one hand this setting could cause that selfish individuals survive because they do not lose fitness because of making costs by helping. But on the long run, no other individuals would be willing to help such selfish individuals. So after all a helpful attitude (low, even negative helper threshold) might be the evolved strategy: If I help now, I will get help later. These investigations might explain situations where an individual makes cost by helping, which occur often in the real world.

7 Conclusions

This paper addressed the discrepancy between cooperating and diverse populations on the one side and the evolutionary principle which seemingly failed to explain this on the other side. The issue was addressed by means of computer simulations, in which cooperation mechanisms, as have been postulated by evolutionary biologists, were implemented to see if this could lead to the evolution of a cooperative and diverse population. The results of the experiments showed that the availability of (innate) cooperation mechanisms in a given population can indeed, in some situations, lead to the evolution of a population of cooperating individuals. Moreover, if there is a lot of cooperation between individuals, groups tend to become more divers and individuals more specialized. In sum, the research establishes by means of computer simulation that cooperation mechanisms can indeed serve as an explanation for the evolution of a diverse population with specialized individuals.

Furthermore the results indicate that the possibility to establish cooperation is one of the fundamental aspects of evolution. I agree with what Nowak [14] states: cooperation should be added as a third fundamental principle of evolution beside mutation and selection.

8 Acknowledgements

I want to thank Ida Sprinkhuizen-Kuyper and Iris van Rooij for being my supervisors. I thank Iris for being my broadening, explorative operator in the ‘evolution’ of my research. I thank Ida for being my narrowing, focussing operator in the ‘evolution’ of my research and helping me with the technical realization.

Also thanks to Oliver Kretschmann for brainstorming with me and giving me feedback.

References

[1] C. Anderson and N.R. Franks. Teams in animal societies. *Behavioral Ecology*, 12:534–540, 2001.

[2] C. Anderson and N.R. Franks. Teamwork in animals, robots and humans. *Advances in the Study of Behavior*, 33:1–48, 2004.

[3] T. Balch. Learning roles: Behavioral diversity in robot teams. Technical report, AAAI, 1997.

[4] T. Balch. Behavioral diversity as multiagent cooperation. SPIE ’99 Workshop on Multiagent Systems, Boston, 1999.

[5] D.I. Bolnick, R. Svanbck, J.A. Fordyce, L.H. Yang, J. M. Davis, C.D. Hulsey, and M.L. Forister. The ecology of individuals: Incidence and implications of individual specialization. *The American Naturalist*, 161:1–28, 2003.

[6] S. Bowles and H. Gintis. The evolution of strong reciprocity: Cooperation in heterogeneous populations. *Theoretical Population Biology*, 65:17–28, 2004.

[7] A.E. Eiben and J.E. Smith. *Introduction to evolutionary computing*. Berlin: Springer, revised edition, 2007.

[8] D. J. Futuyma and G. Moreno. The evolution of specialization. *Annual Review of Ecology and Systematics*, 19:207–233, 1988.

[9] T. Haynes, S. Sen, D. Schoenefeld, and R. Wainwright. Evolving a team. Technical report, AAAI, 1995.

[10] J. Henrich and R. Boyd. Division of labor, economic specialization and the evolution of social stratification. *Papers on Economics and Evolution*, #0720:1–42, 2007.

[11] L. Lehmann and L. Keller. The evolution of cooperation and altruism - A general framework and a classification of models. *Evolution Biology*, 19:1365–1376, 2006.

[12] S. Luke, L. Panait, G. Balan, S. Paus, Z. Skolicki, E. Popovici, K. Sullivan, J. Harrison, J. Bassett, R. Hubley, and A. Chircop. ECJ 18 - A Java-based evolutionary computation research system. Website, 2009. Available online at: http://www.cs.gmu.edu/∼eclab/projects/ecj; visited on Juli 6th 2009.

[13] D. Nettle. The evolution of personality variation in humans and other animals. *American Psychologist*, 61:622–631, 2006.

[14] M. A. Nowak. Five rules for the evolution of cooperation. *Science*, 314:1560–1563, 2006.

[15] M. Quinn, L. Smith, G. Mayley, and P. Husbands. Evolving teamwork and role-allocation with real robots. *Artificial Life*, 8:302–311, 2002.

[16] J.L. Sachs, U.G. Mueller, T.P. Wilcox, and Bull. J.J. The evolution of cooperation. *The Quarterly Review of Biology*, 79:136–160, 2004.

[17] Y. Suzuki and T. Arita. A comprehensive evaluation of the methods for evolving a cooperative team. *Artificial Life and Robotics*, 10:157–161, 2006.

Appendix

A Manipulation of the variable ‘ExtraCosts’

In a pilot study I manipulated the variable ‘ExtraCosts’. This variable indicates how much working alone on a task costs extra. Several values between 1.1 and 1.7 (with pr = 0.5) were tested. The effect is that when there are less extra costs there is less working together, especially on same tasks there is no working together at all. When there are high extra costs there is much working together, as working alone is very expensive (See Figure A-1). In the experiments the variable ExtraCosts is set to 1.3.

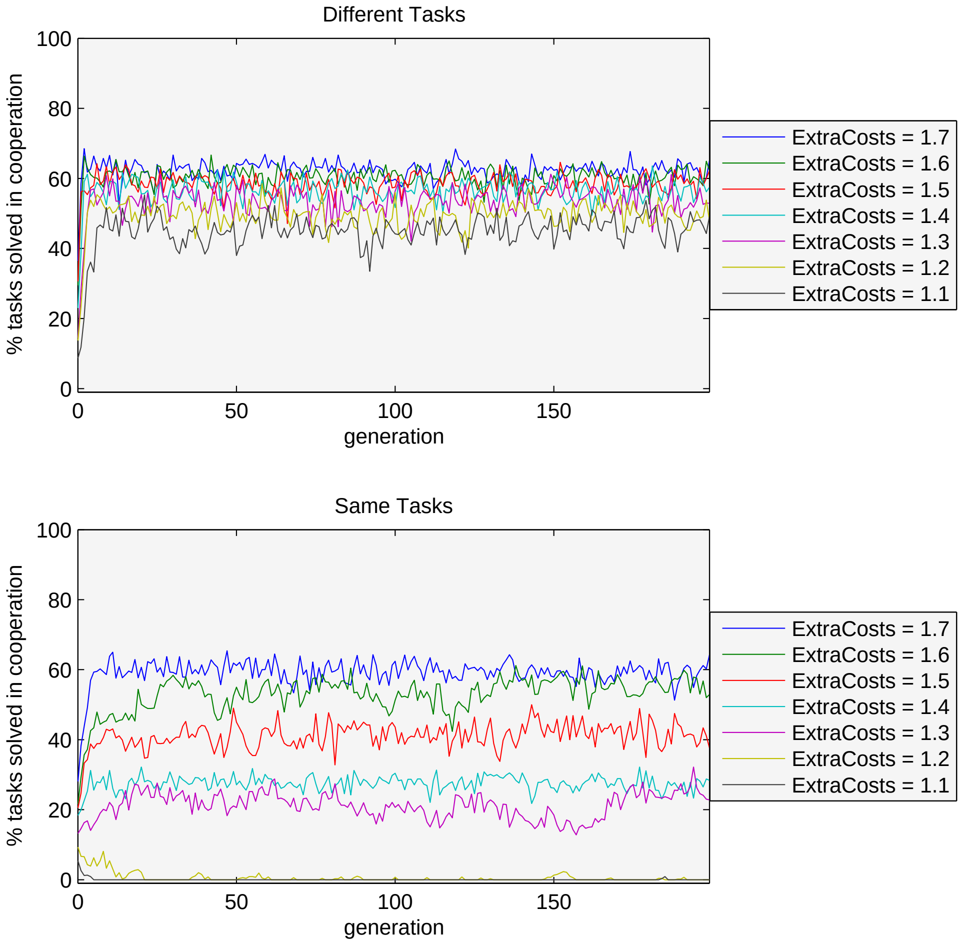

Figure A-1: Influence of the proportion of reward that the helper gets on the amount of cooperation. Top: Different tasks, bottom: Same tasks

B Baseline and maximum values for average fitness

Different average fitness values could be observed for the setting without cooperation mechanisms, depending on the task types. This difference in baseline values is due to the difference in the reward as can be seen in the computations:

Simple tasks (one subtask): Assuming an individual who is average in all skills (skill level = 0.33), working alone:

$$Fitnesscontribution = Reward - (Cost * 1.3) = 10 - ((1/0.33) * 1.3) = 6.1$$

For 10 tasks this yields the baseline fitness value of 61.

Different & Same tasks (two subtasks): Assuming an individual who is average in all skills, working alone:

$Fitnesscontribution = Reward - (Cost * 1.3) = 20 - ((1/0.33 + 1/0.33) * 1.3) = 12.2$

For 10 tasks this yields the baseline fitness value of 122.

Maximal obtainable average fitness value for Different and Same tasks can be computed as follows. Assuming two specialized individuals (skill level max. = 0.8), working together:

$Fitnesscontribution = Reward - Cost$
$Fitnesscontribution_{ind} = (20 - 5) - (1/0.8) = 13.75$
$Fitnesscontribution_{help} = (20 - 15) - (1/0.8) = 3.75$
$Fitnesscontribution_{ind+helper} = 20 - (1/0.8 + 1/0.8) = 17.5$

For 10 tasks this yields the maximal fitness value of 175.

C Gain of working together for different and same tasks

There is generally more gain of working together for different tasks than for same tasks. This can be illustrated with the following examples. Assume the first individual has the skills (0.2, 0.2, 0.6) and the 3 potential helpers have the skills (0.7, 0.2, 0.1), (0.2, 0.7, 0.1) and (0.1, 0.2, 0.7).

Different tasks Assume the task requires the skills A and C. For working alone, this yields:

$Fitnesscontribution = Reward - (Cost * 1.3) = 20 - ((1/0.2 + 1/0.6) * 1.3) = 11.33$

The most optimal choice for working together is with the first potential helper. This yields the following fitness:

$Fitnesscontribution = Reward - Cost$
$Fitnesscontribution_{ind} = (20 - 5) - (1/0.6) = 13.33$
$Fitnesscontribution_{help} = (20 - 15) - (1/0.7) = 3.57$

Fitness gain for the individual when working together: $13.33 - 11.33 = +2$
So solving the task together is the optimal choice.

Same tasks Assume the task requires two times skill C. For working alone, this yields:

$Fitnesscontribution = Reward - (Cost * 1.3) = 20 - ((1/0.6 + 1/0.6) * 1.3) = 15.67$

The most optimal choice for working together is with the last potential helper. This yields the following fitness:

$Fitnesscontribution = Reward - Cost$
$Fitnesscontribution_{ind} = (20 - 5) - (1/0.6) = 13.33$
$Fitnesscontribution_{help} = (20 - 15) - (1/0.7) = 3.57$

Fitness gain for the individual when working together: $13.33 - 15.67 = -2.33$
So solving the task alone is the optimal choice.

D Parameters

Several parameters were used in my evolutionary algorithm. As an overview I list all parameter settings here (see Table D-1 and Table D-2).

From how many individuals a helper can be chosen	NR_HELPERS	3
When expected pay-off ¿ threshold a helper helps	HELPER_THRESHOLD	0
How much working alone costs extra	EXTRA_COSTS	1.3
Populationsize	pop.subpop.0.size	100
Number of generations	generations	200
Length of genome	genome-size	3
Minimum value of genes	min-gene	0.1
Maximum value of genes	max-gene	1
Mutation probability	mutation-prob	0.33
Mutation type	mutation-type	gauss
Standard deviation of Gaussian	mutation-stdev	0.05
Selection method	pipe.source.0	Tournament Selection
Tournament size	select.tournament.size	3

Table D-1: Fixed parameters

Task type (what proportion of tasks is simple, different or same)	PROP_SIMPLE	100	0	0
	PROP_DIFFERENT	0	100	0
	PROP_SAME	0	0	100
Proportion of the reward the helper gets	PROP_REWARD	0	0.25	0.5

Table D-2: Manipulated parameters