Robots that specialize and make exchanges

Domenico Parisi
Institute of Cognitive Sciences and Technologies
National Research Council
Rome, Italy
domenico.parisi@istc.cnr.it

Giovanni Sirio Carmantini
School of Mathematics and Computing
University of Plymouth
Plymouth, United Kingdom
sirio.mail@gmail.com

***Abstract*— Specialization and exchange are two important specifically human adaptations that are at the origin of much of the complexity of human social life and of human societies. The paper describes simple simulated robots that evolve in environments containing either two types of food or both food and tools and tries to establish in which environments specialization emerges, what is the relation between exchange and specialization and what their advantages are.**

Keywords—Specialization; exchange; evolutionary robotics

I. Introduction

Unlike nonhuman animals, human beings obtain most of the goods they need not from nature but from other human beings through the exchanges of goods. X gives one good to Y and, in exchange, Y gives a different good to X. But, to be really useful, the exchange of goods requires specialization in the production of goods. X must specialize in producing one good and Y must specialize in producing another good, and then X and Y can proceed to an exchange. Specialization and the exchange of goods are such an important component of the social life of human beings and of the organization of their societies that, if we want to construct robots that not only have the external morphology of human beings (humanoid robots) but actually behave like human beings (human robots), our robots must be able to exchange their goods and must develop specialization in the production of goods. Much work has been done on economic behaviours in agent-based systems (see, for example, [1]) and in specialization in swarm robotics context [2,3,4], but less work on robots with a body and a behaviour which is controlled by an artificial neural network (see [5,6] for works on specialization in evolving neuro-robots, see [7] for a robotic model of free market). In this paper we describe neuro-robots that evolve in a variety of environments and we try to determine which environments favour the emergence of specialization and the exchange of goods and what their consequences are.

II. Methods

A. Robot characteristics

The robots are simulated Khepera robots. The robots live in an environment which contains two different types of food tokens which the robots can distinguish on the basis of their colour. The two types of food tokens contain two different types of energy and the robots need both types of energy to remain alive and reproduce. Their body has a store for each type of energy, each type of energy is consumed by a fixed quantity at each time step, and if the energy contained in one or the other bodily store reaches the zero level, the robot dies. Since a robot generates an offspring at regular intervals, the robots that live more generate more offspring and therefore leave more copies of their genes to the next generation. (All the robots die at some maximum age which is the same for all the robots.) The robots' behaviour is controlled by a neural network with one layer of sensory units, one layer of internal units and one layer of motor units (Fig. 1). The input layer is composed by four sensory neurons, two neurons for each type of tokens, which are simulated eyes that detect the tokens located within a certain distance from the robot. The higher the activation of the neurons, the nearer the tokens. Like our eyes each couple of sensory neurons produce an overall field of view of 180 degrees, composed by a right and a left part, with a zone of overlap in the center. The hidden layer is composed by four neurons, with an associated bias and a sigmoid activation function. The output layer includes two motor neurons which encode the speed of the two wheels that allow the robots to displace themselves in the environment. The motor neurons also have a bias and a sigmoid activation function. The connection weights and bias values of a robot's

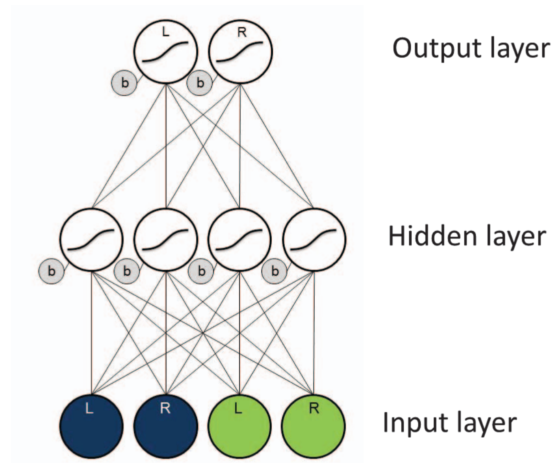

Figure 1. Architecture of the neural network controlling the behaviour of the robot. The input layer consists of two couples of sensory neurons, encoding left and right FOV for each type of tokens; the hidden layer is composed by four neurons and their biases; the output layer is composed by two motor neurons and their biases, controlling the speed of the left and right wheels.

neural network are inherited from the robot's parent (reproduction is nonsexual) with the addition of some random noise which, in some cases, can make the offspring robot better than the parent robot at eating the food tokens. A food token is captured by a robot when the robot reaches (touches) the food token. When a robot captures a food token, it puts the food token into an external store and the robot eats the food tokens contained in its external store when its bodily energy for that type of token reaches a sufficiently low level.

B. *Simulations characteristics*

The initial population includes 200 robots but then the size of the population changes and the population can become extinct because no robot is able to live enough to be able to generate an offspring. The entire simulation lasts for 30000 time-steps (input/output cycles of the robots' neural network) and it is repeated 20 times starting with randomly generated connection weights and biases for the robots of the initial generation. We evolve the robots in four different environments which vary as a function of three variables:

1 - The environment can contain two types of food tokens or only one type of food tokens but also tool tokens. A tool token is a token which cannot be eaten but can be used to triple the energy value of a food token.

2 - The tokens can be randomly distributed in the entire environment or the tokens of one type can only be found in one zone of the environment and the tokens of the other type can only be found in another, separate, zone of the environment.

3 - The robots can or cannot make exchanges.

The robots that cannot make exchanges can only eat the food tokens they are able to find in the environment. If the robots can make exchanges, in each cycle the program randomly selects two robots with complementary external stores and an exchange takes place. "Complementary" means that the external store of one robot contains more food tokens of one type and the external store of the other robot contains more food tokens of the other type. "Making an exchange" means that one token of one type is transferred from the external store of one robot to the external store of the other robot, and vice versa, according to the present state of the two robots' external stores. Once the two robots have made an exchange, they have to wait 20 cycles to participate to a new exchange with the same or another robot. (The number of 20 cycles between two exchanges for the same robot has been arbitrarily decided by us but this number should be varied to determine how it can influence the robots' behaviour.) An exchange is cost-free and leaves no space for the robots to decide whether to participate to an exchange or not and what are the terms of the exchange. This is a simplification but the purpose of this study is not to explore the emergence of exchange and the behaviour of making exchanges, but rather to specify under what environmental conditions exchange is more adaptive and what the relation between exchange and the emergence of specialization is. (We return to these issues in the conclusions.)

C. *Specialization test*

To test if some robots are specialized in capturing one type of tokens and other robots are specialized in capturing the other type of tokens, we put a robot in an "experimental environment" that contains only one type of tokens and we count the number of tokens that the robot is able to capture in a fixed period of 200 cycles repeated for 10 trials. Then we repeat the procedure using the other type of tokens. If a robot captures more tokens of type a than tokens of type b, we say that the robot is specialized in capturing type a tokens. The Specialization Index is defined as

$$S(t_a, t_b) = \frac{t_a - t_b}{t_a + t_b} \tag{1}$$

where t_x is the total number of type x tokens collected by the robot tested. The Specialization Index is a number between -1 and 1 where a value toward the extremes of the range means a specialization in one of the types of food tokens and a value near 0 means no specialization.

III. RESULTS

A. *Food tokens*

1) Random distribution of tokens and no exchange

In the first environment the food tokens are randomly distributed in the entire environment and the robots cannot exchange the food tokens of one type with the food tokens of the other type with other robots. In this environment all the populations of robots become extinct before the end of the simulation (30000 cycles) (Fig. 2). The robots need both type of food tokens to remain alive but, since they can only count on themselves, they often find themselves with an external store which lacks one or the other of the two types of food, and they die.

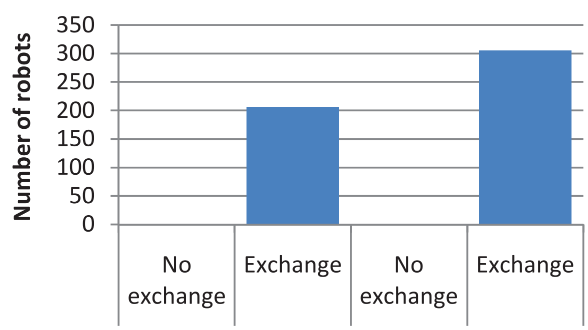

Figure 2. Population size at the end of the simulations, averaged over the 20 replications of each simulation, for the robots living in the environment with two types of food tokens.

2) Random distribution of tokens plus exchange

In the second environment the food tokens are also randomly distributed in the entire environment but the robots can now proceed to exchange one token of one type with one token of the other type with another robot. This helps the robots and at the end of the simulation (after 30000 cycles) 7 out of the 20 populations are still in existence (Fig. 2). So the possibility to make exchanges represents an advantage for these robots. However, in this environment the possibility to make exchanges does not lead to specialization. Fig. 3 shows the frequency of the Specialization Index for this environment discretized in five ranges and computed for the entire set of the 20 replications of the simulation. These robots do not appear to be specialized in capturing one or the other of the two food tokens even if they can exchange the food tokens with other robots. This notwithstanding, the possibility to make exchanges helps the robots to better survive in the environment.

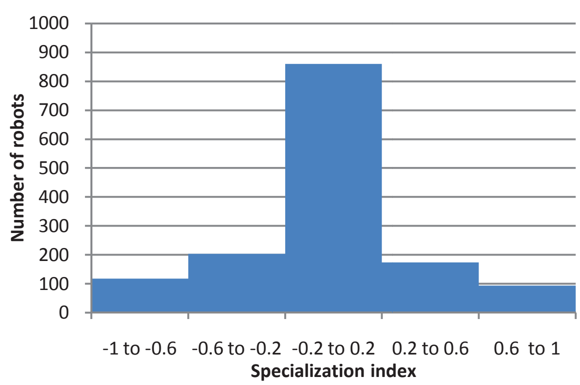

Figure 3. Specialization Index frequency for the environment with randomly distributed food tokens and no exchange. A value near to 1 means specialization in one type of food tokens, whereas a value near to -1 means specialization in the other type of food tokens.

3) Zonal distribution of tokens and no exchange

In this new environment the food tokens are not randomly distributed in the entire environment but the food tokens of one type are in one zone and the food tokens of the other type are in another zone of the environment, and the two zones are separated by an empty space. This is a more difficult environment because the robots still need to eat both types of food tokens to remain alive but to capture both types of food tokens they must traverse the empty space between the two zones and they may die during their travel between the two zones. In fact, if the robots do not have the possibility to exchange the food tokens with other robots, all the populations become extinct even earlier than the populations living in an environment in which the two types of food tokens are randomly distributed in the entire environment.

4) Zonal distribution of tokens plus exchange

Now we add the exchange of the two types of food tokens. Two robots are randomly selected from the entire population of robots and, if their external stores are complementary, the robots exchange one type of food token with the other type of food token. This has two consequences. The first consequence is that 5 out of 20 populations succeed at avoiding extinction (Fig 2). This, again, shows that the exchange of goods is a useful addition to the adaptive pattern of the robots.

The second consequence is that we find specialization. We test individual robots in the "experimental environment" and we find that most robots are specialized. Some robots capture more food tokens of one type than food tokens of the other type and other robots do the opposite. Since they can exchange their food tokens with other robots, the robots have evolved a behaviour which allows them to avoid traversing the empty space between the two food zones. Some robots live in one food zone and they specialize in capturing the food tokens that they find in that zone while other robots do the same for the other zone. Then they exchange the food tokens in which they are specialized with other robots which are specialized in the other type of tokens. Fig. 4 shows the Specialization Index frequency for this population of robots.

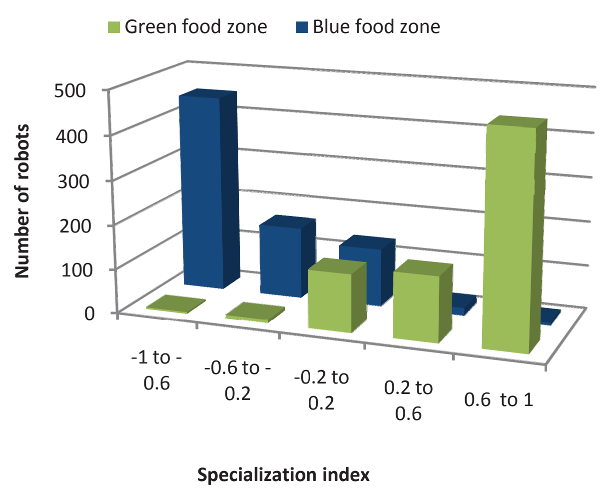

Figure 4. Specialization Index frequency for the two sub-populations (green and blue) which live in the environment with zonal distribution of the two types of food tokens and the possibility to make exchanges.

B. Food tokens and tool tokens

These new robots live in an environment which contains only one type of food tokens and to remain alive they need only the energy contained in these food tokens. However their environment also contains tool tokens. A tool token is identical to a food token except that it has a different colour and when a robot reaches a food token it captures the food token and puts the tool token into its external store together with the food tokens. Clearly, the tool tokens cannot be eaten but they are equally useful to the robots because they increase by three times the quantity of energy contained in a food token. In other words, when a robot eats one of the food tokens contained in its external store and it also has a tool token in its external store, while the food token by itself contains only one unit of energy, by using the tool token the robots can extract three units of energy from the food token. (One tool token can be used only once for a single food token and then it must be thrown away.) The new robots live in the same four different environments that we have seen before with food tokens and tool tokens randomly distributed in the entire environment or with food tokens in one zone of the environment and the tool tokens in another zone, and with or without the possibility to exchange a food token for a tool token. Fig. 5 shows the average population size for each of the four environments.

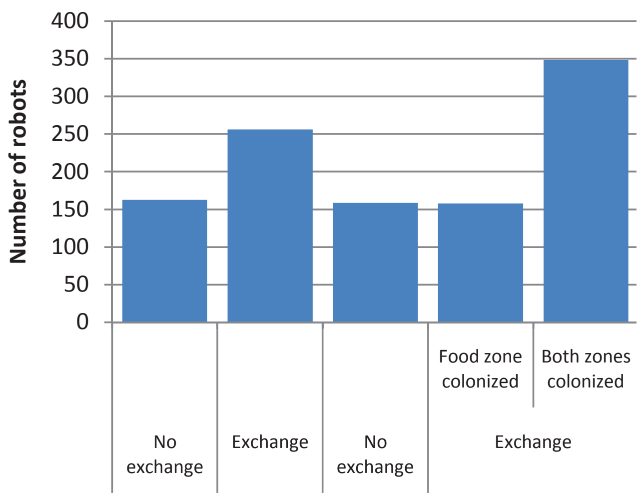

Figure 5. Population size at the end of the simulations, averaged for all 20 replications of each simulation, for the robots living in the environment with food and tool tokens.

1) Random distribution of tokens and no exchange

In the environment in which both the food tokens and the tool tokens are randomly distributed in the entire environment, even if there is no possibility to exchange food for tools, 14 out of 20 populations succeed in surviving until the end of the simulation. Since there is no exchange, the robots do not specialize in capturing food or tools (Fig. 6), although some of the robots exhibit a preference for the food tokens over the tool tokens, which can be explained by considering that the food tokens are needed to survive.

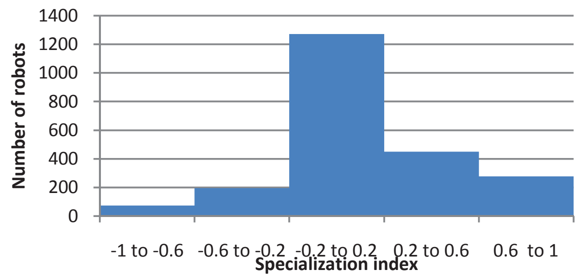

Figure 6. Specialization Index frequency for the population with randomly distributed food and tool tokens and no exchange. A value near to 1 means specialization in food tokens, whereas a value near to -1 means specialization in tool tokens.

2) Random distribution of tokens plus exchange

Now we move to the robots that live in an environment with food tokens and tool tokens randomly distributed but the robots have the possibility to exchange food for tools. In this environment almost all populations (18 out of 20) succeed to survive because, although the robots do not specialize in capturing the food tokens or the tool tokens (Fig. 7), they can optimize their adaptive pattern by having both food tokens and tool tokens in their external store through the exchange of food for tools with other robots. And the average size of the eighteen populations which succeed in surviving is quite large (Fig. 5). Some of the robots show a preference for the food tokens over the tool tokens, as the food tokens are needed to survive while the tool tokens alone don't guarantee survival. This notwithstanding, we do not observe specialization at the population level.

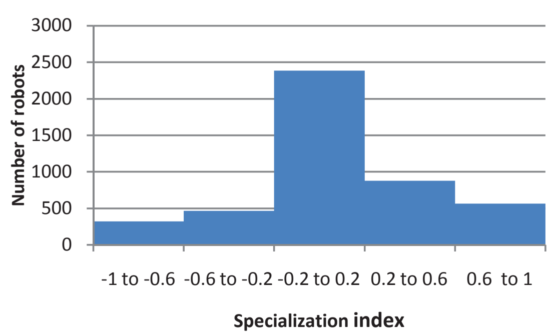

Figure 7. Specialization Index frequency for the population with randomly distributed food and tool tokens and the possibility to make exchanges.

3) Zonal distribution of tokens and no exchange

The robots that live in an environment with a zone containing food tokens and another zone containing tool tokens but that cannot exchange food for tools evolve a different strategy. Since they cannot traverse the empty space which separates the food zone from the tool zone, they all tend to live in the zone containing the food tokens. Although the tool tokens would be very useful to them, they prefer to live in the food zone. And when we test them in an environment containing either only food tokens or only tool tokens, we find that they are specialized in capturing the food tokens (Fig. 8). This an adaptive strategy because all twenty populations succeed in surviving although population size tends to be smaller than for the preceding robots that lived in an environment with randomly distributed food and tool tokens but could exchange food for tools.

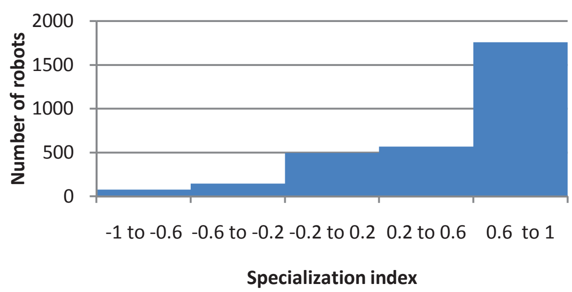

Figure 8. Specialization Index frequency in the population with randomly distributed food and tool tokens and no exchange.

4) Zonal distribution of tokens plus exchange

And, finally, we arrive to the robots that live in the environment with a food zone and a tool zone separated by an empty space but the robots can exchange food for tools. All the twenty populations succeed in surviving in this environment but what is interesting that different populations adopt different strategies. More than half populations (12 out of 20) adopt the more reasonable strategy with some robots living in the food zone and specializing in capturing the food tokens and other robots living in the tool zone and specializing in capturing the tool tokens and then proceeding to exchanging food for tools. The robots of the remaining 8 populations live only in the food zone and they survive by eating the food tokens without using the tool tokens. This is a less effective strategy because population size is much greater for the robots that live in the two zones and exchange food for tools than for the robots that adopt the more conservative strategy of living all in the food zone and renouncing to the tool tokens (Fig. 5). The difference between the two strategies is confirmed by our Specialization Index (Fig. 9 and 10). The robots that adopt the first, more advanced, strategy are divided into two groups: the robots that are specialized in capturing the food tokens (farmers) and the robots that are specialized in capturing the tool tokens (artisans). The robots that adopt the second, less advanced, strategy of living in the food zone without using the tool tokens, are all specialized in capturing the food tokens.

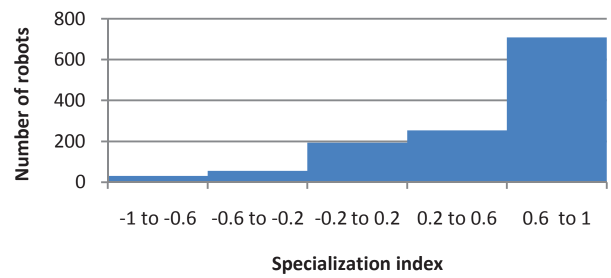

Figure 9. Specialization Index frequency for the populations which live only in the food zone and ignore the tool zone although they have the possibility to exchange food for tools

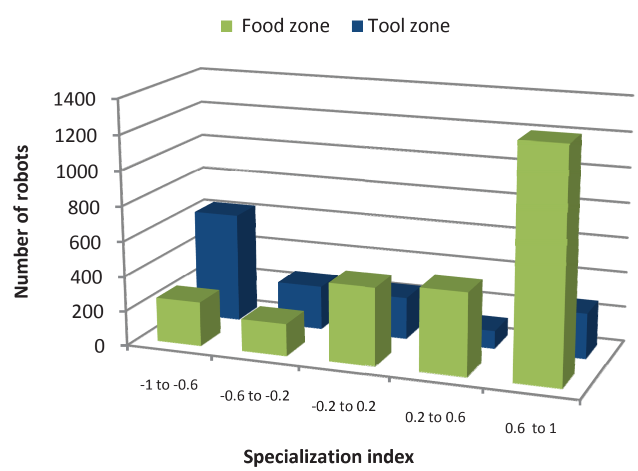

Figure 10. Specialization Index frequency for the robots living in the environment with zonal distribution of food and tools and the possibility to exchange food for tools. Some robots live in the food zone and specialize in collecting (producing) food (farmers, green) while other robots live in the tool zone and collect (construct) tools (artisans, blue) and then the farmers and the artisans exchange food for tools.

IV. Conclusions

Populations of robots that need two different types of food to remain alive and reproduce all become extinct in an environment in which the two types of food are randomly distributed in the entire environment and, even earlier, in an environment in which one type of food is in one zone of the environment and the other type of food is in another zone of the environment. On the contrary, if we allow a robot to exchange one type of food for the other type of food with another robot, many populations succeed in surviving and population size is larger for the robots living in the environment with the two types of food in two separate zones. However, when we measure specialization in individual robots in a laboratory setting containing only one or the other type of food, we observe specialization only for the robots living in the environment in which the two types of food are in different zones. Some robots tend to live in one zone and to collect the food which is found there and then to exchange their food with the food collected by the robots living in the other zone.

If the robots need only one type of food but the environment also contains tools that allow the robots to extract more energy from food, all populations succeed to survive even without the possibility to exchange food for tools because they can adopt the strategy of collecting food without necessarily looking for tools. However, the possibility to make exchanges leads to an increase in population size and the best results are obtained when food is in one zone and tools are in another zone of the environment so that some robots specialize in collecting food (farmers) and other robots specialize in collecting tools (artisans) and then farmer robots exchange food for tools with artisan robots.

Specialization and the exchange of goods are critical human adaptations which have allowed human beings to colonize new and more difficult environments and to develop complex economies. The robots described in this paper make it possible to better understand some basic mechanisms and consequences of specialization and the exchange of goods. Our robots are very simplified with respect to human beings who exchange different types of goods. For our robots there are no costs for exchanging one type of token for another type of token and they automatically proceed to an exchange if one robot has more tokens of one type than tokens of the type and vice versa for the other robot. To explore the role of the costs of making exchanges, one direction of research that we are currently exploring is the emergence of market places. The robots can make exchanges only if they are sufficiently close to one another in physical space. The robot's environment contains a spatial landmark which can be seen by the robots and the robots must evolve the behavior of approaching the landmark and staying near to the landmark (the marketplace) to proceed to making exchanges. This implies that doing exchanges has some costs for the robots in terms of time and, perhaps, fatigue, and we want to see how this additional factor influences the robots' overall behavior.

Another limitation of our robots is that the terms of exchange (one token of one type for one token of the other type) is hardwired by us in the robots. In [8] and [9] we have described more abstract "agents" that autonomously decide whether to proceed to an exchange and what are the terms of their exchange and this important aspect of economic behavior have been explored with other agent-based models. But we want to study these behaviors with our "physical" robots which have a "brain", live in different environments, and possess different adaptive patterns.

A final but very important problem is that in our robots the "farmer" robots have an advantage with respect to the "artisan" robots because the "farmer" robots can survive without tools and, therefore, without exchanging food for tools while the "artisan" robots cannot "eat" their tools and they need to exchange them with food. But, in human history the opposite seems to have happened, with "artisans" living in cities and possessing more goods and "farmers" living in the country and possessing fewer goods. This may be due to the fact that producers of tools can progressively improve their tools so as to make them increasingly desirable to producers of food, while this strategy is much less accessible to producers of food.

REFERENCES

[1] Dascalu, M., Franti, E., & Stefan, G. (1998). Modeling production with artificial societies: the emergence of social structure. *Cellular Automata Research Towards Industry*.

[2] Li, L., Martinoli, A., & Abu-Mostafa, Y. (2002). Emergent specialization in swarm systems. *Intelligent Data Engineering and Automated Learning—IDEAL 2002*, pp. 195-204.

[3] Kernbach, S., Nepomnyashchikh, V. A., Kancheva, T., & Kernbach, O. (2012). Specialization and generalization of robot behaviour in swarm energy foraging.*Mathematical and Computer Modelling of Dynamical Systems*, *18*(1), pp. 131-152.

[4] Brutschy, A., et al. (2012). Costs and benefits of behavioural specialization. *Robotics and Autonomous Systems*.

[5] Nitschke, G. (2007). Neuro-evolution methods for designing emergent specialization. *Advances in Artificial Life*, pp. 1120-1130.

[6] Eiben, A., Nitschke, G., & Schut, M. (2007). Collective specialization for evolutionary design of a multi-robot system. *Swarm Robotics*, pp. 189-205.

[7] Dias, M. B., & Stentz, A. (2000, July). A free market architecture for distributed control of a multirobot system. In *6th International Conference on Intelligent Autonomous Systems (IAS-6)*,pp. 115-122

[8] Cecconi, F. and Parisi, D. (2007), *"Asymmetric pricing: an agent based model"*, In Proceedings of the IASTED International Conference on Modelling and Simulation (MS 2007), February 13 - 16, Montreal, CANADA. , pp. 380-385.

[9] Delre, S.A. Parisi, D. (2007), Information and cooperation in a simulated labour market: a computational model for the evolution of workers and firms. In M Salzano, D. Colander (Eds.) *Complexity Hints for Economic Policy*. New York, Springer,