A Real-World Ecosystem featuring several Robot Species

Tony Belpaeme and Andreas Birk
Vrije Universiteit Brussel, Artificial Intelligence Laboratory
Pleinlaan 2, 1050 Brussels, Belgium
{tony,cyrano}@arti.vub.ac.be

Abstract

This paper presents an artificial ecosystem embedded in the real world. The ecosystem hosts different robotic species linked together in their search for "food" in form of electrical energy. The paper deals with four species: the so-called "moles", the "competitors", the "mouse", and the "head". The differences between species are quite substantial. For example the "moles" have simple sensing capabilities whereas the "mouse" is equipped with vision. The "competitors" are a kind of parasite whereas the "head" depends on cooperation. These differences lead us to the notion of the so-called intelligence in terms of features. We present its relation to self-preservation and what follows is an additional support of the viewpoint that for an animat its environment has to be taken into account.

1 Introduction

Following the "artificial life route to artificial intelligence" [19] the notion of intelligence can be based on the concept of behavior. Steels defines behavior as intelligent "if it maximizes preservation of the system in its environment"[20]. It follows that the environment itself plays an important role when investigating animats. This opinion is also expressed and motivated in e.g. [9, 30].

The "natural" way of making a environment complex is to use several different animats in an ecosystem-like setting. But most of the Alife research featuring ecosystems is based on simulations (e.g. [3, 6, 29, 12, 13, 27]). In that case, perception and effector-control of the agent are not embedded in the real-world. The various disadvantages of this are commonly known and agreed upon. A simulated complex environment is a kind of contradiction in itself. To generate a complex environment in a simulation immense modeling is needed. Using the real world, the researcher is spared this overhead as the "world is its own best model" [7].

When it comes to robotic ecosystems, or simply multi-agents-systems based on robots, usually settings with a single "species" are used. This means, the robots used in one setting are all more or less the same (e.g. [28, 8, 18, 14]). The few labs with several different robots commonly work with a "zoo-mentality", i.e. the animats are "held" separately in their "cages". The opportunities of interaction in an ecosystem are not exploited.

At the VUB AI-lab, we use an artificial ecosystem featuring animats from several different species. The differences are substantial as they include much more than just size and weight. First of all, the animats in this ecosystem have very different "hardware-features" as some have e.g. a vision module, others not. In addition, they differ in their "complexity of accessing food". Some can directly access a charging-station for re-filling their batteries, others depend on interactions with the first ones.

The rest of the paper is organized as follows. Section 2 describes the basic set-up of the VUB ecosystem including a charging-station, simple robots, and the so-called competitors. In section 3 the notion of feature-intelligence is introduced. Its relation to self-preservation is discussed and its influence on the conceptualizing of further species is presented. In section 4 the so-called "mouse" is described. This robot features vision which is used to detect the charging station and competitors. Section 5 introduces a somewhat unusual autonomous robot as it is immobile. This so-called "head" consists of a camera on a pan-tilt-unit with quite some vision capabilities. It "sells" information to other robots to prevent itself from starving. Section 6 concludes the paper and discusses future work.

2 Autonomous Recharging

"Food" in form of electrical energy is the basic force merging the individual animats into our ecosystem. It is a crucial component to keep the animats *viable* [17, 15, 1]. Most of the animats are mobile[1] and powered by secondary batteries. These batteries can be re-filled in a charging station. Picture 1 shows the basic set-up as presented and explained in detail in [21, 16]:

The basic ecosystem consist of

[1] One animat, the so-called "head", is not mobile and therefore has to rely on other robots. It is described later in this paper.

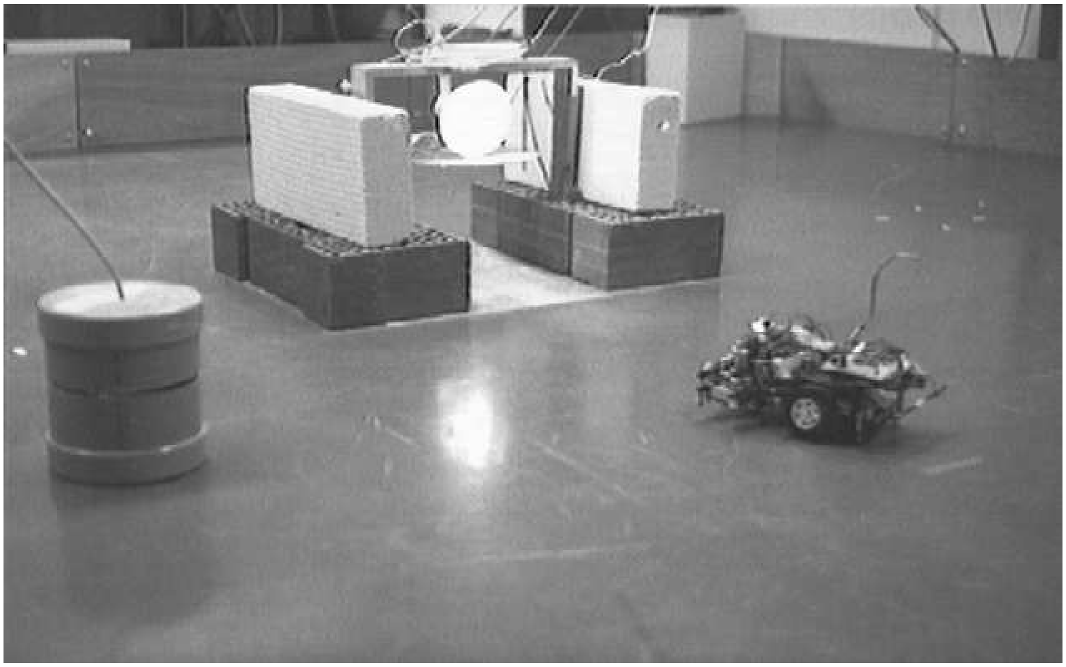

Figure 1: The basic ecosystem: the charging station (middle), a competitor (left) and a robot (right).

the charging-station with a bright white light on top of it

simple mobile robots , the so-called "moles"[2] which can do photo-taxis towards the charging-station and towards

the competitors , these are boxes housing lamps connected to the same global energy-source as the charging-station. They can be dimmed by other inhabitants of the ecosystem by being pushed.

The basic ecosystem already provides various possibilities for Alife research. The interested reader is referred to e.g. [24, 4, 23, 22].

3 Intelligence in terms of features vs self-preservation

As mentioned in the introduction, intelligence can be linked to the capability of self-preservation of a system. In this section we will have a look on another way of "defining" intelligence and compare the two approaches.

This second way is more related to "classic" AI and takes account of the fact that intelligence is often identified with having certain capabilities. For example Turing identifies in his famous test, intelligence with the capability to take part in a conversation based on written natural language. Other such capabilities are e.g. vision, learning, planning, and so on (loosely speaking a list of all the subfields of classic AI). We will denote the fact of having such a capability as *feature*. The *feature-intelligence* of an agent is simply the sum of all features he possesses. Usually, some kind of order is implied on features. Some are seen as contributing more to "intelligence" as others. But we will not bother about that here.

At a first glance, feature-intelligence and self-preservation are linked in a straightforward way: one would say that the more features the better the agent is suited to survive. But this is not necessarily true. Features have a negative impact on self-preservation as well. As mentioned above, energy is a crucial factor in keeping viable. Unfortunately, additional features require additional hardware in terms of sensors and effectors as well as computation power. Therefore, they lead to an increase in the energy consumption.

What follows is an additional reasoning for the philosophy that the environment has to be taken into account. In simple environments a simple set of features will do. An "overkill", like e.g. having "cognitive" robots with high-level planning facilities in the above described basic ecosystem, is seen as good in respect to feature-intelligence. But it is bad in terms of self-preservation as the planning will require much more computational

[2]The names for this and following species should not be taken too literal. The name "moles" attributes to the fact, that these robots have no vision. We use these names for reasons of convenience, not to imply or suggest any deeper relation with the biological versions.

power as the boards for our simple robots can provide. So, bigger boards are needed, faster CPUs, more memory, and so on. The robots get bigger and the energy-consumption goes up. Especially, the consumption will increase much more than high-level planning can help to find "food" (in the very simple basic setting).

Of course, additional features can be very useful depending on the environment. The more complex and/or changing the environment is, the more and better features are needed to achieve adaptivity. But features should be treated in a realistic way. Often, hardware factors like computational power, energy consumption, size, and so on, are simply neglected as "minor technical details". But they strongly influence self-preservation in real-world settings.

Let us illustrate this in an example. Using a small robot, a radiolink, and a "invisible" number-crunching Cray next door should not be sold as a design of a simple small animat. Instead, the Cray must be made "visible", i.e., its energy-consumption has to be taken into account in the ecosystem. Imagine how much "food" for thousands of our "moles"[3] would be available by replacing a single "dinosaur"-Cray.

4 The feature vision and the "mouse"

Among the most crucial features for any animat are sensing-capabilities. Unfortunately, most of todays sensors are very specialized. Embedding adaptation on the sensor-level is somehow hard. Of course, many applications of learning to sensors are known. But they usually deal with increasing accuracy and robustness. Changing the functionality, like e.g. turning an active IR-sensor from an obstacle-avoiding device into an edge-detector, is infeasible for most sensors.

This does not hold in respect to vision. The same combination of a camera, a digitizer, and a computing-device can be used for obstacle-avoidance as well as edge-detection. Therefore, we chose vision as sensor-feature to play an important role in the "diversification-process" in the ecosystem.

The first new species equipped with vision is the so-called "mouse" (figure 2). The mouse is a kind of enhancement of the "moles" by adding a camera at the front of a robot. At the moment, we cheat a little bit by using radio-transmission of the camera-pictures to a PC-host which does the processing. But we are investigating the possibility to use a Phytec TI320C50 DSP-board with a piggyback frame-grabber (the whole hardware is smaller than a cigarette-pack) to do the complete job on the robots.

[3] A "mole" consumes roughly 5 W.

4.1 The visual analysis

For the visual analysis we closely stick to the active vision paradigm several properties which makes it the ideal solution for the visual performance we are aiming at. The purposive and task-oriented design — meaning that only what is needed is computed —, the use of cues and attentional mechanisms and the lack of elaborate internal representations reduce the computational complexity. Reducing the need for computational power is most important when using low-cost, portable hardware. Other properties are the tolerance to errors, the dependence on recognition rather than reconstruction, the roots in biological vision systems and the link to behavior based systems.

Several parallel-working modules handle the visual perception, each module handles a certain cue or task. The input to each module is the unaltered camera image or this image fed through some low-level analysis (typically some Sobel-edge detection and some basic smoothing). The architecture, with the in (simulated) parallel working modules, is shown in figure 3.

4.2 Seeing the competitors

The competitor-module has as tasks the detecting of competitors in the image and returning their position, a distance-estimate and their state (active or stunned). This problem can be tackled in a number of ways, e.g. by doing template matching or by using some neural network. We however use a very straightforward solution. All competitors are dark colored and a simple thresholding of the image reveals all dark patches. All these patches are checked for their aspect ratio. If the aspect ratio is within a certain margin the same as the aspect ratio of a competitor, it is considered being one. This method works very fast and is very satisfactory for our purpose; the few times that it does not correctly recognize a competitor (this happens when two competitors appear as one dark patch in the image) don't pose a problem. So recognizing the competitors is reduced to finding dark patches having a specific aspect ratio. Next, the state of each competitor is checked. Each competitor houses a lamp emiting modulated light —this is used by the "moles" to do photo-taxis using simple sensors— if the lamp is on, the competitor is active, if it is off the competitor is stunned (and the robot no longer homes in on the competitor). The competitor module checks for this light, and marks every competitor as active or not.

Estimating the distance to a competitor is simple. The competitors are cylindrical and look the same from every angle, so the distance to the observer is reversely proportional to their size in the image (see (1) where a and b are constants depending on the aspect ratio of a competitor and the units in which one would like to express the distance).

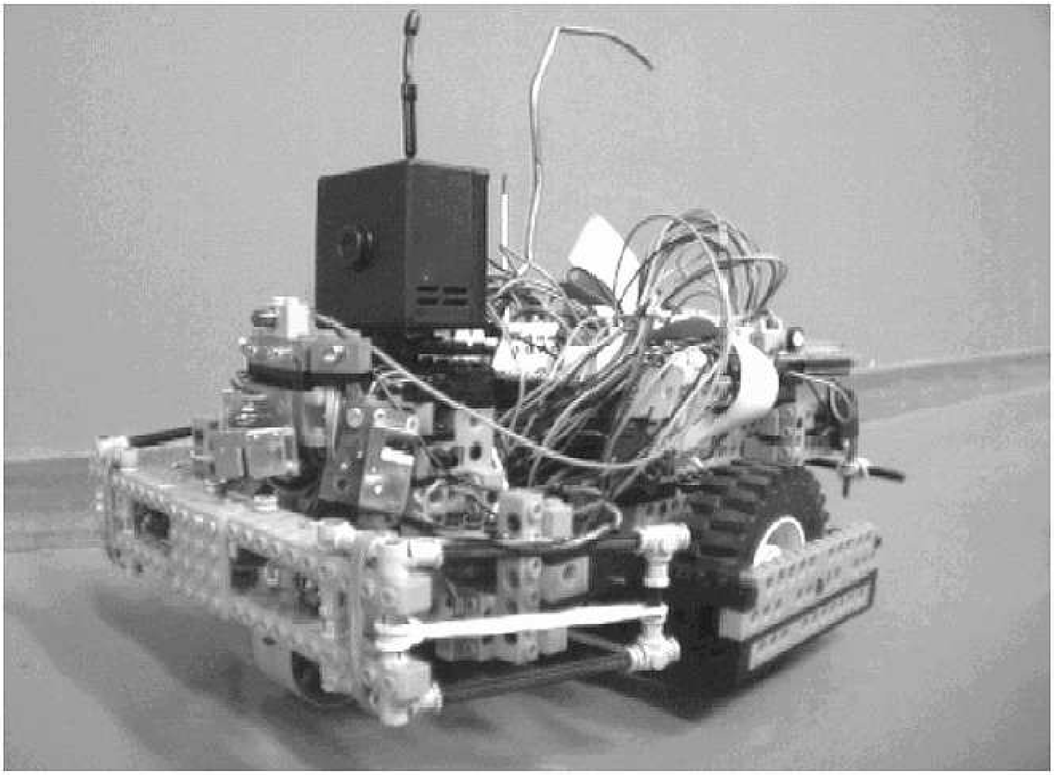

Figure 2: The robot "mouse" with its camera at the front.

$$distance = \frac{1}{2}\left(\frac{a}{width} + \frac{b}{height}\right) \tag{1}$$

The module builds a radial depth map containing all parasites, their position in the image, their distance to the robot and their status (active or not). All this information is available to the behavior system of the robot, but of course not all this information is needed. The "align on competitor"-behavior only needs the position of the closest competitor alive; while the obstacle avoidance module needs the position of all competitors, no matter whether they are alive or not.

4.3 Seeing the charging station

The charging station has a bright white light, which is used by the "moles" to home in on the charging station using only their two light sensors. Since the white light is the most prominent feature of the charging station, it is also used by the visual system to recognize the charging station. The module thresholds the incoming image for light colors and calculates the centroid. To avoid confusion caused by the overhead lights and their reflections on the shiny floor, the centroid is calculated in a restricted region in the middle of the image.

Recognizing the charging station with other algorithms (template matching, Hough-transforms,...) is extremely hard. This is mainly due to the charging station's irregular shape, to the fact that is can be viewed from an infinite combination of distances and angles and to the ever changing reflections of light in the ecosystem.

The floor of the ecosystem is flat and the camera viewpoint is always on the same level (the camera cannot tilt), so objects will appear at the same height in the image. This is called the *ground-plane constraint* and can be used to estimate the distance to objects, since farther objects will appear higher in the image. Calculating the distance to the charging station was however not easy. Using the position of the charging station in the image, or its size seemed an obvious choice; but in practice the position and the size are very hard to determine because of the same problems we experienced during recognizing the charging station. In a last attempt it was fitted with black horizontal stripe, which is clearly visible and allows for a distance-measure (the height of the stripe in the image is inversely proportional to the distance).

4.4 Seeing obstacles

When the *ground-plane constraint* is respected, it is surprisingly easy to do obstacle avoidance; provided the floor of the environment is smooth (such as the uniform

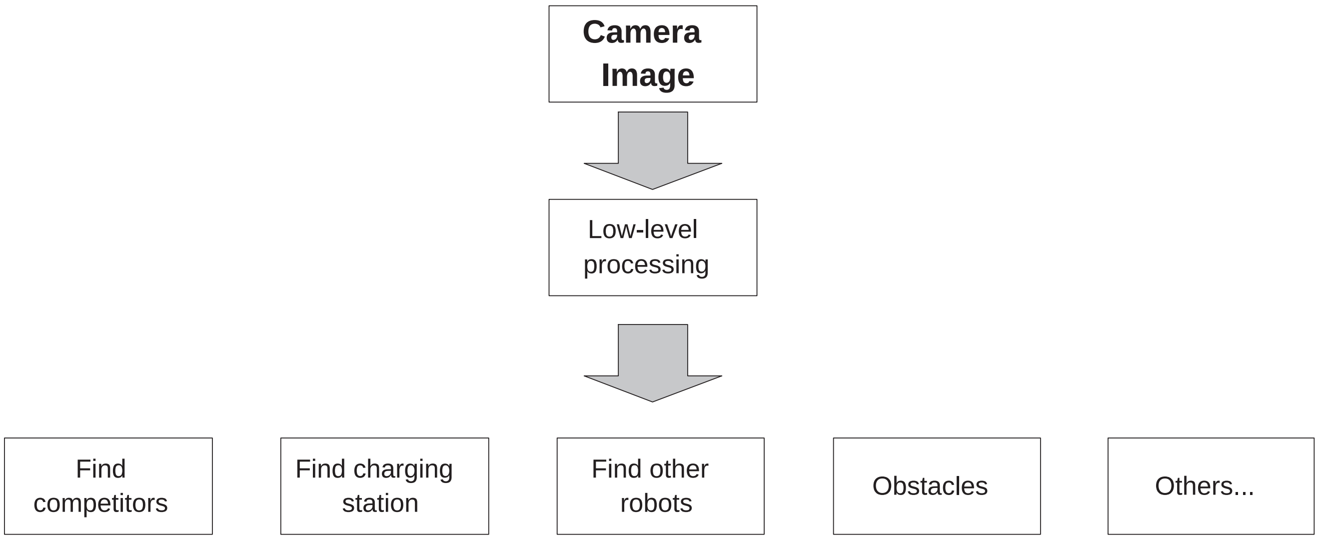

Figure 3: The architecture of the vision system. Several modules each asynchronously handle a cue or task.

colored floor of the VUB AI-lab ecosystem). The lower part of the image always shows the floor, when doing edge detection this part of the image will show no or hardly any edges. Unless an obstacle is close enough to the robot to enter this lower image part. So walls, other robots and objects will be detected as soon as they come within a predefined distance to the robot. As a result the obstacle avoidance module returns the direction furthest away from all obstacles, the behavior system for the control of robot can use this to steer clear from all obstacles. This method has also been used by Horswill in Polly [10].

4.5 Seeing other robots

In some experiments (e.g. [25]) a robot needs to see other robots. Since robots are the only thing moving in the ecosystem, the "Find other robots"-module only needs to look for motion in the image not caused by the observer itself (this implies that robots have to move in order to be seen). One could use optical flow analysis for this, but since this is a computationally expensive way to handle the problem, we use difference images. This has one drawback over the optical flow-approach: the observer can not move while looking for other robots. This is however acceptable in our current experiments.

The "Find other robots"-module also monitors the ego-motion of the robot. This can be useful when the robot gets stuck. The robot bodies don't carry any wheel encoders, and the ego-motion information can be compared to the last motor actuation; if these don't match a retract-behavior should be stimulated.

5 Forced to cooperate: the "head"

The so-called "head" is the most recent "inhabitant" of the VUB ecosystem. It is a very unconventional autonomous robot as it is not mobile. But though being forced to stay at one place, far away from the "food-source", it can preserve itself from starving. As we will see, it can provide useful information to other (mobile) animats. It does not do this for "free", but for a certain portion of the benefit in terms of energy. So, it is kind of "fed" by other robots in exchange for providing information.

The head consists of a camera mounted on a pan-tilt-unit (figure 6) and is equipped with quite some computing power. As we want it to have strong vision-capabilities the hardware needed is not feasible to be carried around by small or medium-sized robots. Following the reasoning in section 3, and due to the interesting aspects involved in this constraint, we decided to conceptualize the head as immobile animat.

5.1 The capabilities of the "head"

Initially the head makes a 180 degree scan of the environment, building a radial depth map containing all competitors. After that it switches into a watch-dog mode, where it just randomly looks around the ecosystem. When its attentional mechanisms (which trigger on unusual motion) pick up something interesting, the focus is placed on that particular region. Since the only things moving in the ecosystem are the robots, the head will pick out a robot and start tracking it. The head will try to predict its path and will utter a warning (over the

Figure 4: A competitor as seen by the robot. A border is placed around the competitor, meaning that is has been recognized as active. To the right the charging station can be seen.

Figure 5: The charging station, correctly recognized by the robot at a distance of 2m.

radio-link) if the robot comes close to a relevant object. Kuniyoshi et al. [11] built a system doing the same, but instead used a stereo head and more powerful hardware.

For the scanning of the ecosystem, the head rotates

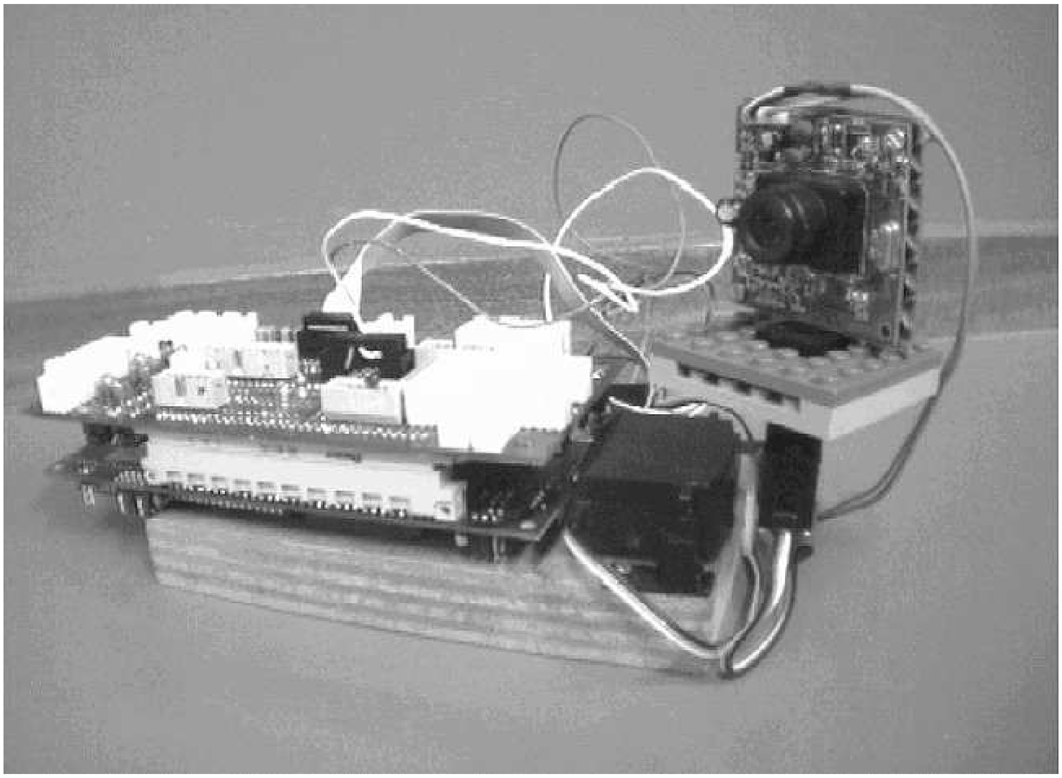

Figure 6: The “head”, an inhabitant of the ecosystem forced to cooperate.

its pan-motor step-wise from left to right and analyses each image frame, meanwhile building the radial depth map of all things relevant. In the watch dog mode the head observes a part of the ecosystem for a few seconds, when nothing of interest is detected it performs a saccade and observes another part of the ecosystem. During the saccade the head is unable to see anything, due to its egomotion.

Finally, the tracking is done using difference images. As soon as the robot is about to move out of the image, the camera is repositioned to get to robot back into the middle of the image. We call this *saccadic tracking*. The head also continuously compares the position of the robot with the position of the competitors in the radial detpth map, if a robot comes too close to a competitor it is warned using the synchronous radio link every robot carries.

The architecture of the head's visual analysis is basically the same as in figure 3, consisting of several modules running in parallel. All analysis and actions happen real-time.

6 Conclusion and Future Work

We presented a real-world ecosystem including several robots species. The species introduced in some detail, as they are presented for the first time in this paper, were the “mouse” and the “head”.

The mouse is a kind of enhanced version of the “moles”. The moles are simple robots which where described previously in several publications. In addition to simple sensors like the moles, the mouse is equipped with vision. This allows it, among other capabilities, to recognize the charging-station and the competitors. In doing so, two improvements over the moles are achieved. First, the mouse can see the charging-station and the competitors over much longer ranges. Second, it can estimate their distances.

At the moment, the vision processing of the mouse is done via a radio-connection on a host-pc. This contradicts somewhat our philosophy described in section 3 that the technical requirements of so-called features must not be ignored. But as the processes involved are rather computational cheap, we belief that they can be embedded on a DSP-board on the mouse itself. We hope we can prove this in future work.

Further future work with the mouse must deal with the question of what is gained in terms of self-preservation with vision. As mentioned above, the vision improves perception of the charging-station and the competitors. But it leads as well to an increase in the energy-consumption. The impacts of these two opposing forces have to be investigated in appropriate long-term experiments.

The "head" is the second new inhabitant of our ecosystem. It consists of a camera on a pan-tilt unit and is immobile. The immobility follows from our credo of obeying technical constraints. In this paper we describe how the head can help other animats by providing information: its is capable of tracking robots and issuing helpful warnings. As it cannot get to the charging-station itself, it relies on "selling" these informations to other robots.

The mechanisms for this cooperation have not yet been investigated by us. So, following questions belong to the most important future work in respect to the head: How can a ratio of receiving energy for giving information be determined? How can cheating and free-rides be prevented? How can the cooperation evolve or be learned? And many more.

7 Acknowledgments

Many thanks to Werner van Belle for writing parts of the "heads" software (environment scan and distance prediction to competitors). The robotic agents group of the VUB AI-lab is partially financed by the Belgian Federal government FKFO project on emergent functionality (NFWO contract nr. G.0014.95) and the IUAP project (nr. 20) CONSTRUCT.

References

[1] Ross Ashby. *Design for a brain.* Chapman and Hall, London, 1952.

[2] Dana Ballard, Animate Vision. *Artificial Intelligence*, **48** (1991) 57-86.

[3] Hugues Bersini. Reinforcement learning for homeostatic endogenous variables. In *From Animals to Animats 3. Proc. of the Third International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1994.

[4] Andreas Birk. Learning to survive. In *Fifth European Workshop on Learning Robots, Bari*, 1996.

[5] Andrew Blake and Alan Yuille, *Active Vision,* MIT Press, Cambridge, Massachusetts, 1992.

[6] Bruce Blumberg. Action-selection in hamsterdam: Lessons from ethology. In *From Animals to Animats 3. Proc. of the Third International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1994.

[7] Rodney Brooks. Intelligence without reason. In *Proc. of IJCAI-91.* Morgan Kaufmann, San Mateo, 1991.

[8] Robert Ghanea-Hercock and David P. Barnes. An evolved fuzzy reactive control system for cooperating autonomous robots. In *From Animals to Animats 4. Proc. of the Fourth International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1996.

[9] Ian Horswill. Characterizing adaption by constraint. In *Toward a Practice of Autonomous Systems, Proceedings of the First European Conference on Artificial Life.* The MIT Press, Cambridge, 1992.

[10] Ian Horswill. Polly: A Vision-Based Artificial Agent, In *Proceedings AAAI-93*, Washington, US, 1983.

[11] Y. Kuniyoshi, M. Ishii, S. Rougeaux, N. Kita, S. Sakane and M. Kakikura, Vision-Based Behaviors for Multi-Robots Cooperation, In *IEEE Int. Conf. on Intelligent Robots and Systems*, 1994.

[12] Long-Ji Lin. Self-improving reactive agents: Case studies of reinforcement learning frameworks. In *From Animals to Animats. Proc. of the First International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1990.

[13] Pattie Maes. A bottom-up mechanism for behavior selection in an artificial creature. In *From Animals to Animats. Proc. of the First International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1990.

[14] Maja J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In *From Animals to Animats 2. Proc. of the Second International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1993.

[15] D. McFarland and A. Houston. *Quantitative Ethology: the state-space approach.* Pitman Books, London, 1981.

[16] David McFarland. Towards robot cooperation. In Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson, editors, *From Animals to Animats 3. Proc. of the Third International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1994.

[17] Jean-Arcady Meyer and Agnes Guillot. Simulation of adaptive behavior in animats: Review and prospect. In *From Animals to Animats. Proc. of the First International Conference on Simulation of Adaptive Behavior.* The MIT Press/Bradford Books, Cambridge, 1991.

[18] Alexandros Moukas and Gillian Hayes. Synthetic robotic language acquisition by observation. In *From Animals to Animats 4. Proc. of the Fourth International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1996.

[19] L. Steels and R. Brooks, editors. *The "artificial life" route to "artificial intelligence".Building situated embodied agents.* Lawrence Erlbaum Associates, New Haven, 1993.

[20] Luc Steels. The artificial life roots of artificial intelligence. *Artificial Life Journal, Vol 1,1*, 1994.

[21] Luc Steels. A case study in the behavior-oriented design of autonomous agents. In Dave Cliff, Philip Husbands, Jean-Arcady Meyer and Stewart W. Wilson, editors, *From Animals to Animats 3. Proc. of the Third International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1994.

[22] Luc Steels. Discovering the competitors. *Journal of Adaptive Behavior, 4(2)*, 1996.

[23] Luc Steels. A selectionist mechanism for autonomous behavior acquisition. *Journal of Robotics and Autonomous Systems 16*, 1996.

[24] Luc Steels and David McFarland. *Cooperative Robots: A case Study in Animal Robotics*. The MIT Press/Bradford Books, Cambridge, 1994.

[25] Luc Steels and Paul Vogt, Grounding adaptive language games in robotic agents. Submitted to the 4th European Conference on Artificial Life, Brighton, UK, 1997.

[26] Michael Swain and Markus Stricker, editors, Promising Directions in Active Vision, *International Journal of Computer Vision,* **11** (1995) 109-126.

[27] Toby Tyrrell and John E. W. Mayhew. Computer simulation of an animal environment. In *From Animals to Animats. Proc. of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1990.

[28] Barry Brian Werger and Maja J. Mataric. Robotic "food" chains: Externalization of state and program for minimal-agent foraging. In *From Animals to Animats 4. Proc. of the Fourth International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1996.

[29] Gregory M. Werner. Using second order neural connections for motivation of behavioral choices. In *From Animals to Animats 3. Proc. of the Third International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1994.

[30] S.W. Wilson. The animat path to AI. In *From Animals to Animats. Proc. of the First International Conference on Simulation of Adaptive Behavior*. The MIT Press/Bradford Books, Cambridge, 1991.