Proceedings of the 2007 IEEE Symposium on
Artificial Life (CI-ALife 2007)

Evolution of Tool Use Behavior

Boris Schäfer
Dept. of Computer and Information Science
Linköping University
58 183 Linköping, Sweden
borsc@ida.liu.se

Nicklas Bergfeldt,
María José Riveiro Carballa
and Tom Ziemke
School of Humanities and Informatics
University of Skövde
Box 408, 541 28 Skövde, Sweden
tom.ziemke@his.se

***Abstract*— This paper focuses on the capability of artificial evolution to produce tool use behaviors of different complexity in simulated robotic agents and in the absence of learning or other lifetime methods. The results show by example that tool use behaviors of different complexity can evolve and do not necessarily rely on reasoning abilities.**

I. Introduction

Evolutionary robotics takes advantage of artificial evolution in order to create autonomous agents. The advantage of using artificial evolution in the design process is the reduction of human design bias [1]. One important feature of robotic autonomy is the ability to identify environmental objects as potential tools that can be used by an agent in order to solve its task. However, in evolutionary robotics the evolution of tool use behavior has received hardly any attention so far. This might be due to the assumption that higher-level tool use behavior requires too complex sequences of fine-granular sub-behaviors. Artificial evolution develops successful strategies by means of generalisation and does not incorporate specific one-time events of short period experience. Robotic tool use has so far either been pre-programmed, as in industrial robotics, or developmental lifetime approaches have been applied, such as learning or object recognition algorithms, see e.g. [2, 3]. Pre-programmed behaviors provide the precision that is needed in industrial robots but not the flexibility that is expected of truly autonomous agents. Developmental approaches on the other hand make use of sophisticated hardware and software imposing great demands on lifetime resources and yielding relatively complex agents in contrast to the idea of simplicity as promoted by evolutionary robotics. Some researchers go another way by introduction of ambient intelligence [22] or artifact intelligence [21]. However, these approaches distribute the autonomy over several elements as opposed to a single autonomous agent. The work described in this paper shows by example that artificial evolution by itself can yield tool use behavior of different complexity. That means, tool use of different complexity is achievable without the agent learning or reasoning about it during its lifetime. While it is true that only the combination with lifetime methods will finally yield useful engineering results this paper describes initial investigations of the limits of artificial evolution regarding tool use complexity. Another motivaton for investigating tool use it the proposal of a new research direction for evolutionary robotics that has been made in [16], the proposal of investigating the possibility of an agent actively adapting the environment to its needs instead of only adapting the agent itself to the environment as this is usually done in evolutionary robotics experiments. This has earlier been pointed out in [18] from a human perspective. While [16] described initial experiments based on stigmergy as observed in insects, in this work detached tool use, i.e. the use of environmental objects as tools, is considered as another approach following this direction. There has been related work within the artificial life community. In contrast to our investigation they used crudely simplified grid-world simulations. For example, in [19] stigmergy was investigated from an information-theoretic perspective and [20] investigated social learning of tool use behavior. The latter proposed social learning as an effective alternative to imitation learning. In contrast to our work [20] applied learning methods and outlined the combination with an evolutionary dimension for future work.

II. Background

There are no commonly agreed definitions of tools and tool use because their qualities depend on context rather than inherent properties. Obviously, there exists a circular relationship between tools and tool use, i.e. the one is usually used to define the other. Therefore tool and tool use definitions are either too narrow (e.g. human-centered) or too general to fit all possible scenarios. A thorough overview of this problem and various approaches towards definitions can be found in [17]. In this paper a tool is an environmental, persistant, physical object that becomes a tool in the moment that it is picked up by a tool-using agent. The successful use of the tool in order to solve the agent's task is then considered tool use. It is important to note that objects that are already part of the agent's body or attached to its body from the start of an experiment (e.g. sensors or grippers) are not considered as tools within this paper, thus a tool must exist as a detached environmental object prior to use. This is important because the complexity of identifying an object as a potential tool, approaching it and using it in a manner that is beneficial to the user's task seems to demand a higher level of deductive analysis, reasoning and planning than using parts of one's own morphology. Consumable or non-persistant items such as food or energy are also not considered as tools within the context of this paper. It is also noteworthy that random success must be excluded from being claimed to be tool use. There needs to be a significant difference in task efficiency when tools are used and it must be possible to replicate such results. A common view on tool use behavior is that complex tool use behavior can only be learned during an agent's lifetime as it requires capabilities to reason about concepts of the world or exploration facilities such as understanding cause-effect loops, see e.g. [4]. While some researchers claim that traits captured by natural evolution might be supportive in the acquisition of tool use behavior [5] it is questionable whether natural evolution alone would be able to capture more complex or fine-granular tool use behavior. This seems to be obvious considering that the direction of evolutionary search is dependent on the life-time experience of former generations, an experience that is abstracted and generalized by the evolutionary process. However, the limits of purely evolved tool use as defined in this paper have not been tested in robotic simulation so far. We consider this study as an initial test and encourage other researchers to take part in the detection of these limits. Beside the contribution to the theoretical body of knowledge regarding the relations between tool use, evolution and development this might also guide future engineering approaches that try to combine artificial evolution and developmental approaches, thereby potentially off-loading an agent's long-term and general conceptualizing efforts of the world to artificial evolution.

III. Experiments

This paper focuses on initial investigations of limits of evolved tool use behavior and presents only those setups and results of a series of experiments that are relevant to this question. The target of the experiments was to show by example that evolved tool use behavior of different complexity is possible, even if no lifetime learning is applied. The experiments can be considered as initial steps towards a better understanding of the interplay of evolutionary and learning mechanisms adapting tool use behavior. Obviously, there are countless possible variations of the experiments and many different tools and tool use scenarios that can be thought of. However, our initial experiments were designed with simplicity in mind in order to achieve both comparable results for analysis and examples of evolved tool use behavior that indicate that artificial evolution alone could indeed produce more complex behavior than commonly expected.

A. Experimental setup: The experiments have been performed using the simulator YAKS that allows the simulation of multiple Khepera robots [7]. Simulations never perfectly reflect the real world due to different factors, such as hidden implementation bias, the availability of perfectly precise sensors and other factors such as environmental influences that the researcher might not be aware of (e.g. dynamic light conditions) (cf. e.g. [1, 8, 9]). Therefore simulated experiments usually need to be validated in real-world experiments. Choosing YAKS is regarded as a compromise given the costs and efforts that real-world experiments have because it implements sampled measurements taken on real Khepera robots [10]. The experimental framework was inspired by experiments of Búason (cf. [12, 10, 11]) who investigated pursuit problems in co-evolutionary experiments. In his experiments a predator agent had to catch an evasive prey agent. In the experiments described in this paper the co-evolutionary aspect was not part of the investigation, but the pursuit problem was considered a suitable task for an initial setup in order to test tool use behavior as it resembles the survival task of predators and preys in nature. Basing the framework on Búason's work allowed the verification of the basic implementation in YAKS. The environment is a square with a side-length of 468 mm (cf. Figure 1) which allows the camera with a range of 500 mm and a view angle of 360 degrees to cover the whole environment. The diameter of the robotic agents is 55 mm and their infrared proximity sensors have a range of 40 mm.

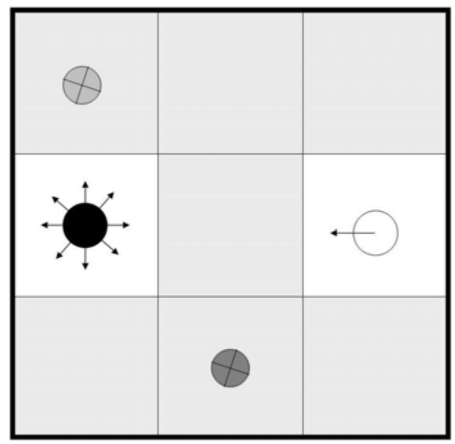

Fig. 1. Experimental environment. The thick lines represent the walls (that can be detected by the agents, visually and by the infrared proximity sensors). The black and white circles represent the predator and the prey, respectively. Arrows indicate possible start directions. The white zones are the fixed starting zones for the predator (left) and prey (right). The small circles with crosses in represent the tools that can be placed on one of the other zones (gray). Only one tool can be placed per zone (the picture shows one possible placement of the tools). The thin lines show the division into nine logical zones (not visible to the robot; only shown for illustration purposes).

Two different simple tools have been implemented that can only be collected and used by the predator. Both take immediate effect upon collection and are physically present cylinders that can be detected by proximity infrared sensors and cameras. In addition both tools are light sources that can be distinctively detected by two different light sensors that are implemented on the predator. Upon collection of the motor tool, the predator gets a maximum speed equal to the prey's top speed. Prior to collection the predator's maximum speed is only half of the prey's maximum motor speed. Upon collection of the camera tool, the camera of the predator is activated (produces null values prior to the collection of the camera tool). Two simplifications should be noted in particular. First, the perception is simplified by the fact that the two different tools can be distinctively detected by two different light sensors which makes it unnecessary for the agent to classify data coming from one sensor with respect to the two tools. Secondly, the pick up procedure of the tool is regarded as a reflex as proposed in [13] for tool use experiments, thus the pick up procedure was not physically modelled, i.e. if the experiments were to be transferred onto real robots one would need to design these reflexes.

Both agents have six infrared proximity sensors on the front and two on the back. The predator has a light sensor with a resolution of eight digits in the one-tool experiments (i.e. in the experiments where only one of the tools is present in the environment) and an additional one in the two-tool experiments(i.e. in the experiments where both tools are present in the environment). The predator's camera that can be activated upon collection of the camera tool has a resolution of 5 digits, a view range of 500 mm and a view angle of 360 degrees. The network of the prey is a simple feed-forward network that maps the inputs from its eight inputs fed by the eight standard Khepera proximity sensors directly to two motor output nodes. For the predator three different simple modular feed-forward network with two, three and four modules have been used in different experiments, respectively. Fig. 2 presents a schematic overview of the network with three modules. A modular network was chosen due to the assumption that tool use behavior can be regarded as a sequence of sub-behaviors. For example, using the camera to solve the pursuit task implies the sub-behavior of detecting the camera tool, collecting the camera tool, using it to track and finally approach the prey.

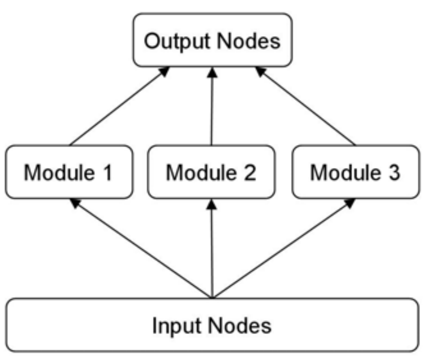

Fig. 2. Schematic view of modular network with three modules.

In the experiments the network weights of the predator have been evolved. Based on the example of natural settings and Búason's experiments [[11], [10]] the fitness function has been designed as a simple time-to-catch function.

$$f(x) = \frac{(step_{max} + 1) - step_{end}}{step_{max} + 1} \text{ , } in\ case\ of\ a\ catch. \quad (1)$$

In (1) $step_{max}$ is the maximum number of time steps for an experiment and $step_{end}$ is the number of time steps that have passed to catch the prey, normalized to a range between 0 and 1. However, 1 is a theoretical maximum that can never be achieved due to the startup distance between the predator and the prey. If there is no catch there is no reward given, i.e. $f(x) = 0$.

Four sets of experiments have been performed that resembled incremental tool use complexities. The first with only the motor tool present, the second with only the camera present, the third with both tools present and the last with both tools present while only the first tool that was picked up had an effect. Thus, when both tools where present combined tool use was possible while tool selection had to be performed in the fourth set of experiments.

In the first two experiments (with one tool present) the lifetime was set to a maximum of 1000 time steps. In the two-tool experiments the maximum lifetime was set to 1500 time-steps. Búason used 500 steps in his experiments but in this work the tool use complication consumes additional time. Each population consists of 100 individuals that are evolved over 100 generations. After each generation elitism and tournament selection was used for the reproduction of the best 20 individuals in combination (thus preventing premature convergence of the populations to a local maximum fitness that might otherwise be achieved by a dominant, sub-optimal strategy). Mutation was applied. These values are inspired by [10,11] in order to achieve some level of comparability (pure elitism was used in his experiments). Different startup scenarios (further called epochs) were used to test the different starting angles of the predator and all permutations of tool placement, concluding in 56 epochs for the experiments with one tool (8 starting angles, 7 tool positions) and 336 epochs for the experiments with two tools (8 starting angles, 42 possible tool position combinations). The prey has been pre-trained in a setting without tools or co-evolution. The prey's fitness function was set to zero if it evaded successfully and to 1000 if it was caught. This allowed the prey to evolve a network that encodes a relatively successful evasion strategy. In the final experiments the network weights of the prey remained static and it was simply used as a target object in a pursuit task. However, in this training a defensive strategy evolved that was not limited by human designer bias and therefore presumably more flexible that a pre-programmed strategy.

More details on the experimental setup can be found in [14].

IV. Results

This section presents only the most significant results from the final generations of the experiments with three modules in the network. The experiments with two and four modules yielded similar results and have been omitted for brevity. The results for all networks can be found in [14], including the progress development over the generations and module use investigations.

A. Success: The one-tool experiments were run for 100 generations, 56 epochs and 30 replications, totaling to 1680 single runs per generation. Table I shows the average results of generation 100 for each experiment rounded to full percentages. In the baseline experiment one tool was present and could be collected by the predator but did not have any effect. In the motor tool experiment the motor tool could be collected and had the effect of doubling the predator's maximum motor speeds. In the camera tool experiments the camera of the predator was activated upon collection.

TABLE I

RESULTS OF ONE-TOOL EXPERIMENTS.

Baseline	Camera tool	Motor tool
20%	55%	59%

Table I shows that the success rate in the final generation was almost tripled in comparison to the baseline experiments with either tool present while displaying a slight advantage for the motor tool experiments. The latter can be explained by the fact that the camera allows the predator to detect the prey from a distance but does not have a positive effect automatically, i.e. using the camera is a more complex sub-task. Moreover, in both experiments with effective tools the tools have been picked up in over 90% of the successful runs of the final generation. In the baseline experiment the tool was only picked up in 34% of the successful runs of the final generation. These results clearly indicate that tool use evolved in both cases. Fig. 3 shows the development of the catch rate over the generations for both one-tool experiments (with the camera tool and the motor tool present in the environment, respectively) and the baseline experiment where the tool had no effect. It is noteworthy that it is not until after approximately 30 generations that the camera tool use displays better success than the baseline experiment while the motor tool use experiment displays better efficiency right from the start. This supports the observation that artificial evolution needs more training to evolve effective camera use due to the greater complexity of exploiting the camera. In addition it can be observed that the catch rate does not seem to stabilize but is still increasing in the final generations of all experiments. It might therefore be possible that the camera catch rate would outrun the motor tool catch rate if the experiments were conducted for more generations, i.e. that after a certain number of generations the complexities of camera tool use might be captured better by artificial evolution. It would therefore be desirable to test this in future work by simply evolving the same experiments over more generations. An initial replication with only 150 generations showed that the catch rate of the camera experiments came already very close to the catch rate of the motor tool experiments in the final generation. However, this could not be validated by more replications due to the time restrictions.

Fig. 4 shows the development of the tool success rates for all one-tool experiments. The tool success rate shows the percentage of success after collection of the respective tool, i.e. the average results for runs where the tool was actually picked up are shown. Also in this case it the development of the camera success rate is delayed by some generations as opposed to the motor tool success rate development. However, both experiments with effective tools display high success

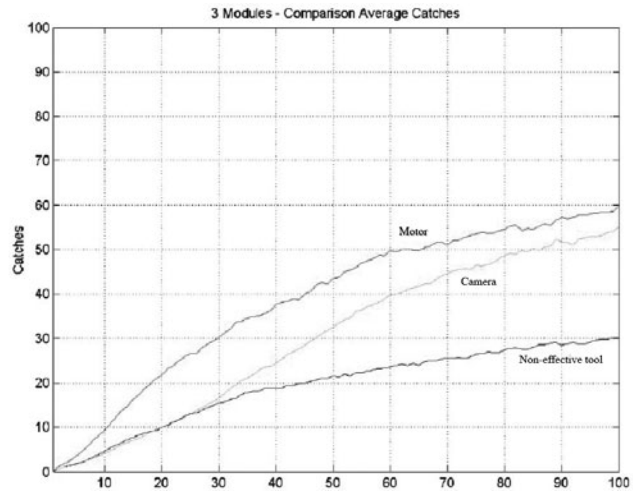

Fig. 3. Catch rate development over generations. Comparison of one-tool experiments with motor tool, camera tool and non-effective tool, respectively. The picture shows average catch rate on the y-axis and the generations on the x-axis. The average results of 30 replications and 56 runs per generation are shown.

rates after few generations and both experiments achieve a success rate of over 90% in the final generations. This clearly indicates that the tools are actually used efficiently once they are collected, i.e. there is clear separation from the random success rates displayed in the baseline experiments. The initial improvement of the baseline experiment success rate can be easily explained because the agent learns to quickly move through the environment in the first generations, thereby increasing the chance of a random catch by covering as much area as possible in time.

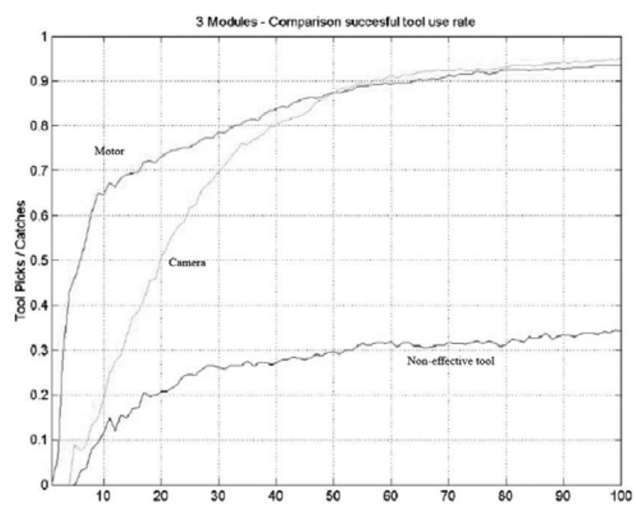

Fig. 4. Tool success rate development over generations. Comparison of one-tool experiments with motor tool, camera tool and non-effective tool, respectively. The picture shows average tool success rate on the y-axis and the generations on the x-axis. The average results of 30 replications and 56 runs per generation are shown.

The two-tool experiments were run for 100 generations, 336 epochs and 30 replications, totaling to 10080 single runs per generation. Table II shows the average results of generation 100 for each experiment, rounded to full percentages. Fig. 5 shows the development of the catch rate over the generations of the two-tool experiments. In the baseline experiment two tools were present and collectible by the predator but did not have any effect. In the tool combination experiment both tools could be collected and effective as described above. In the tool selection experiment both tools were collectible but only the tool that was collected first became effective.

TABLE II
RESULTS OF TWO-TOOL EXPERIMENTS

Baseline	Combination	Selection
50%	80%	72%

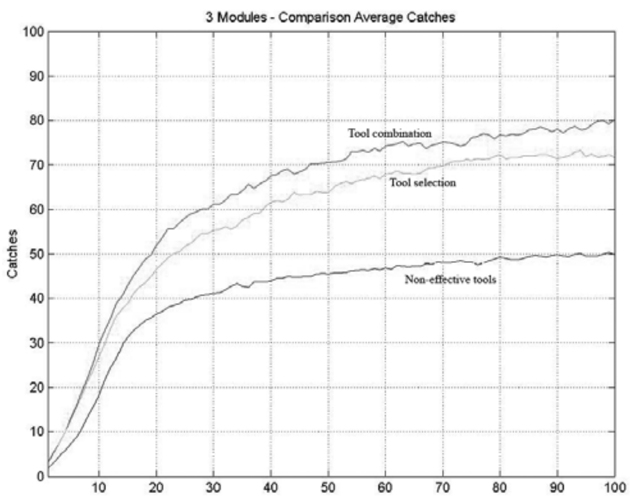

Fig. 5. Catch rate development over generations. Comparison of two-tool experiments with combined tool use, selective tool use and non-effective tools, respectively. The picture shows average catch rate on the y-axis and the generations on the x-axis. The average results of 30 replications and 336 runs per generation are shown.

It is noticeable that the results of the baseline experiment are significantly higher compared to the baseline experiment of the one-tool experiments. This can be explained by two factors. First, the number of epochs was sixfold (due to the possible tool position combinations) and therefore each generation had significantly more lifetime experience. Secondly, uncollected tools were obstacles for the prey. With two tools present, the chance that the prey got trapped by an uncollected tool, and thus easier to catch, was higher in the two-tool experiments. However, both two-tool experiments with effective tools yielded significantly higher success rates (cf. Table I). It is obvious that the success rate of the tool combination experiment is higher than the success rate of the tool selection experiment.

Another noticeable discovery was that in the successful runs of the tool combination experiment the motor tool was picked up in 72% of the cases but the camera tool was only picked up in under 40% of the cases. The latter is nearly in line with the pick up rate of both tools in successful runs. Therefore it can be concluded that the camera tool was rarely picked up alone and there is a clear preference for the motor tool. This can also be observed in Fig. 6 where the development of the camera pick up rate is shown for the different experiments. Thus, a selective tool use behavior evolved even though combined tool use was possible. However, in cases where only the camera was picked up in the tool combination experiment the agent was still successful, explaining the slight advantage over the baseline results. The results for the tool selection experiments show that the camera pick up rate dropped towards a low number of random pick ups after the preference for the motor tool evolves around generation 13. The preference for the motor tool can be explained by the fact that the motor tool is much easier to use. The agent does not need to change its strategy but simply gets an advantage by becoming faster. The camera tool on the other hand is more complex to use as the sensor data needs to be interpreted by the network in order to track the prey. The typical behaviors of the agents are described in Section B.

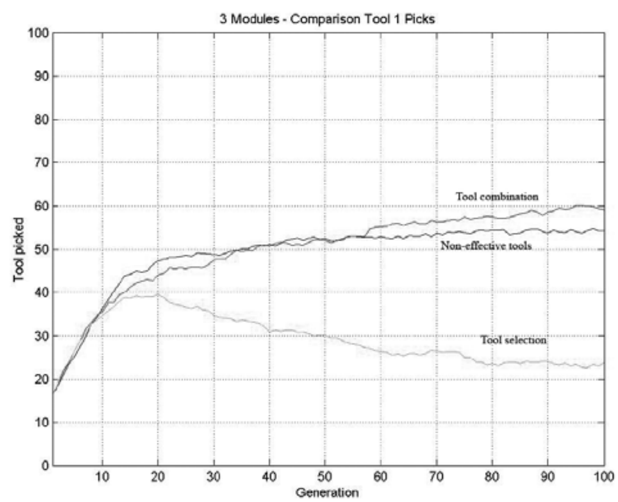

Fig. 6. Camera pick up rate development over generations. Comparison of two-tool experiments with combined tool use, selective tool use and non-effective tools, respectively. The picture shows average camera picks on the y-axis and the generations on the x-axis. The average results of 30 replications and 336 runs per generation are shown.

However, the success rate in the tool combination experiment was higher than the rate yielded in the tool selection experiment, which can be explained by the cases in which both tools were picked up and used successfully. This implies that the artificial evolution of efficient combined tool use is possible. This is also indicated by Fig. 7 that shows the rate at which both tools were picked up in the two-tool experiments. While the baseline results show the rate for the case in which both tools are picked up due to randomness the pick up rate for both tools is significantly higher in the late generations of the tool combination experiments. The importance of picking the camera might increase if remote detection was more important, e.g. in a larger environment. In the tool selection experiment the motor tool was picked up in approximately 83% of the successful runs in the final generations while the camera tool was only picked up in less than 20% of the cases which in the same range with the pick up rate that can be observed in the baseline experiment, i.e. can be explained by accidental picks while moving through the environment. The preference for the motor tool that was already indicated by the results of the tool combination experiment becomes even clearer in the tool selection experiment.

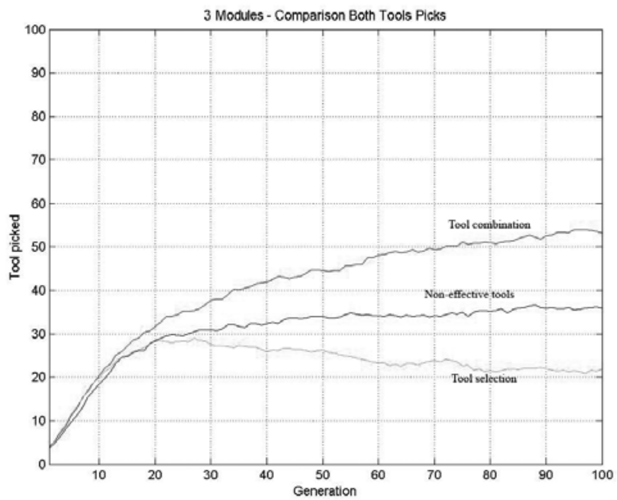

Fig. 7. Pick up rate development for both tools over generations. Comparison of two-tool experiments with combined tool use, selective tool use and non-effective tools, respectively. The picture shows average pick up rate of both tools on the y-axis and the generations on the x-axis. The average results of 30 replications and 336 runs per generation are shown.

Fig. 8 shows that the efficiency improvement of solving the task with tools is not only manifested by the higher catch rates but also by time-to-catch values in the late generations of the tool combination and tool selection experiments that are considerably lower than in the baseline experiments.

B. Behaviors: This section describes the typical behavior of the prey and predators in successful runs of the final generations and applies to all experiments. The typical behavior of the prey was to stay close to its starting point and spin round. This increases the chance to detect an approaching predator in a timely manner because it offers the best sensory coverage relying exclusively on the near-range infrared proximity sensors. When a predator approaches the prey fled in the opposite direction and continued its spinning strategy in another place if not followed by the predator. If it was followed by the predator the typical situation was a chain of small circles in which the agents followed each other. If the prey could not escape the predator this finally ended with a catch in a situation where the prey got obstructed by approaching a wall or a tool in most cases. The strategy that the predator evolved without effective tools present in the environment was to move through the environment in large semi-circles,

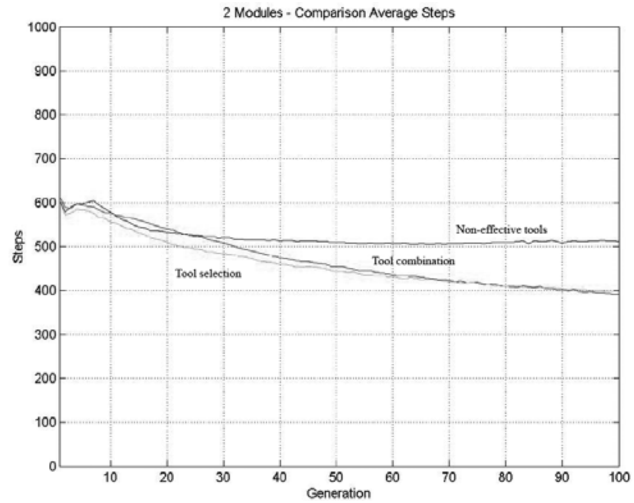

Fig. 8. Time step development for successful runs over generations. Comparison of two-tool experiments with combined tool use, selective tool use and non-effective tools, respectively. The picture shows average steps to a successful catch on the y-axis and the generations on the x-axis. The average results of 30 replications and 336 runs per generation are shown.

changing the direction only when approaching a wall. This semi-circular trajectory remained the same during the whole run because one motor was always at full speed while the other controlled the direction. This is a clever strategy in order to reduce the complexity of synchronizing the motor speeds for navigation and most effective, considering that the predator is blind on long distances. This strategy allows the agent to cover as much area as possible and therefore it increases the chance to find the prey within a certain time. When the prey came into the range of the proximity sensors the predator followed the prey which ended up in the trajectory of small circles as described above or, if the prey reacted in a timely manner, the predator lost track of the prey and continued its strategy of a movement in large semi-circles while the prey began spinning around in its new position. When the camera was collected the predator was able to detect the prey from everywhere within the environment. Hence, it could approach the prey in relatively direct manner from a distance. It did not approach in a straight line because one motor still remained at full speed so that it needed to spin around one time in the worst case and approached the prey in a slight semi-circle. When coming close to the prey the strategy was as described before but with the ability to re-track the prey efficiently when it managed to escape out of the close range.

Fig. 9 shows the schematic view of a typical run in which both tools could be used in combination. In line with the general observations described above it can be observed that the predator moves in semi-circles exclusively. However, instead of executing a wide semi-circular movement it approaches the motor tool in a relatively direct manner. Upon collection of the motor tool the predator speeds up (as indicated by the distance between the black oval marks) and continues with the strategy of a wide semi-circular trajectory. Notably, it does not approach the camera, although it is able to detect it in the same way as it detected the motor tool before with its light sensors. It gets near the prey at around 180 time steps but not close enough to detect it. Therefore the predator continues to move in wide semi-circles until it comes close to the camera tool after approximately 230 time steps. Now the predator approaches the camera, i.e. a sharp turn occurs instead of a continuation of the wide-circle strategy. Upon collection of the camera it can now detect the prey from a distance and approach it. The prey is still spinning around close to its starting point B. When the predator is approaching the prey tries to escape but is followed by the predator. Finally, the predator manages to catch the prey in point C.

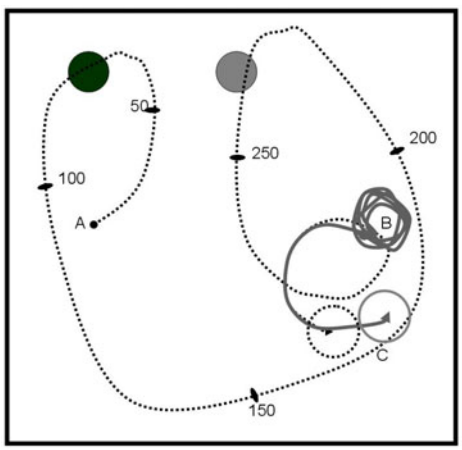

Fig. 9. Combined tool use example. The thick black lines depict the walls of the environment. The filled black circle represents the motor tool in the position it holds from startup of the round until collection (when it is removed from the environment). The filled gray circle represents the camera tool in the same way. The dotted line depicts the trajectory of the predator. Oval black marks demarcate points where the predator was located with the respective number of steps given beside it. The gray line shows the trajectory of the prey. The circle with a dotted frame and the circle with the gray frame show the predator and the prey in the moment of the successful catch, respectively. A and B the starting points of the predator and prey, respectively. The final catch occurs in point C.

C. The role of modularity: The module use investigations in [14] indicate that only up to two modules were frequently used in the majority of the successful runs that were investigated. Fig. 10 shows an example where one module is used nearly all the time except for one time step. We selected an example with three modules in order to show that the third module is, contrary to our expectations, not used at all. This also explains the similarity of the results of the different networks. It might also indicate that the assumption that modularity is beneficent for evolving tool use behavior is not strongly supported. However, we believe instead that benefits of modularity and more complex module use patterns might evolve when the experiments would be run over more generations, possibly allowing more effective straight line movements. The success rates described had still an increasing tendency when the experiments were stopped an more complex networks usually need more training. To support the assumption that modularity plays a significant role even if only two modules are used frequently the experiments could be repeated with a non-modular network.

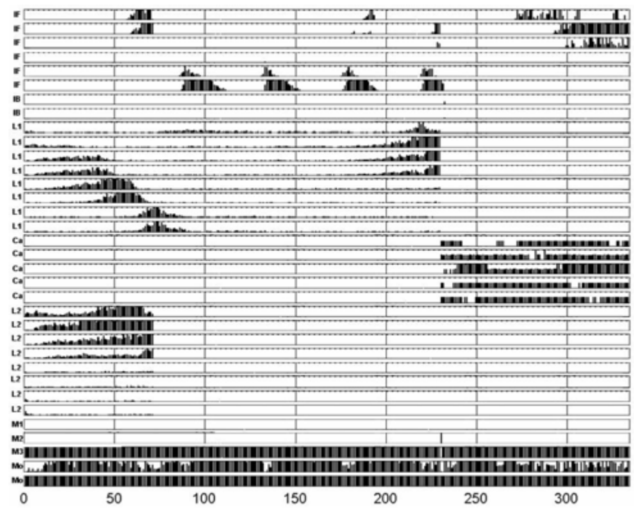

Fig. 10. Module use diagram. This diagram shows the activations of sensor input nodes, actuator output nodes and module activation nodes over time. The time steps are shown along the x-axis and activations along the y-axis (with a minumum of 0 and a maximum of 1). IF and IB refers to front and rear infrared sensor input node activations, respectively. L1 and L2 refers to light sensor 1 and 2 input node activations, respectively. Ca refers to camera input nodes. M1, M2 and M3 refer to activations of modules 1, 2 and 3, respectively. Mo refers to the motor output activation (the lowermost Mo refers to the left motor. All sensor activations are normalized between 0 and 1. Module activation is shown by a black bar whereas no activation is shown by a white bar. Only one module is active at one time. The motor activation is normalized between 0 and 1 where 0.5 means no activation, 1 means full-forward activation and 0 means full backwards activation. The diagram shows the readings for the example run shown in Fig. 9.

V. Discussion and Conclusions

The results of the experiments illustrate that tool use behavior of different complexity can indeed be evolved. These initial experiments have been kept simple. Certainly, more investigations would be desirable, including a much larger variety of setups and parameter settings, different environments, tasks and tools, the implementation of non-distinctive sensors and of pick up procedures, three-dimensional simulation, higher-level tool use and many more. The aim of those experiments should be the establishment of a body of knowledge regarding evolved robotic tool use. Only then one can think of the combination of evolved and learned tool use.

Once we understand evolved tool use behavior and come closer to define the limits of artificial evolution with respect to tool use behavior the next question would be how to combine artificial evolution approaches to autonomous tool use with lifetime approaches such as learning. In order to benefit from this combination we need to find out how to off-load demands on computation power and memory to off-line evolution. This makes it necessary to identify ways that allow to handle long-term and general concepts with artificial evolution while handling the conceptualization of world features that change quickly and unexpected (and can thus not be experienced and captured by evolution) by the application of lifetime methods. In particular we have to solve the question of how to separate these concepts so that lifetime methods ignore the concepts that artificial evolution is meant to take care of. Of course, off-loading an agent's lifetime demands to artificial evolution comes at a cost. Artificial evolution is a time-consuming process and in order to speed this process up simulations are necessary that model the world of the target application sufficiently well in order to produce results that are transferable to real robots.

From a theoretical perspective the experiments are of interest as well. Tool use has been regarded as a hallmark of human-level intelligence, see e.g. [3, 15]. This view has been modified several times in scientific history when examples of animal tool use and even animal tool manipulation and tool manufacture have been observed. As of today, the distinction of human-level and animal intelligence is often claimed to be displayed by higher-level human tool use [6]. One might say that higher level-tool use includes specifially ordered sequences of sub-behaviors, for instance. However, this is quite a blurry distinction and some researchers hypothesized that higher-level tool use in this sense might more be a hallmark of opportunity and necessity than intelligence, a fact that is supported by comparative studies of captive and wild animals where captive animals, that are exposed to more sophisticated objects than can usually be found in the wild, display more complex tool use behaviors.

In this work evolved tool use of different complexity was shown. The paper shows that the limits of evolved tool use complexity are worth investigating. The results of our simulated experiments indicate that they might be pushed towards higher levels of complexity than widely expected. In future a great variety of similar investigations should be carried out in order to clarify those limits.

ACKNOWLEDGMENT

The experiments reported in this paper were carried out by the first author during his time at the University of Skövde [14]. We would like to thank the anonymous reviewers for their valuable comments.

REFERENCES

[1] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence,and Technology of Self-Organizing Machines*, MIT Press, Cambridge and London, 2000, Second printing, 2001.

[2] W. Bluethmann, R. Ambrose, M. Driftler, S. Askew, E. Huber, M. Goza, F. Rehmark, C. Lovchik, and D. Magrunder, "Robonaut: A robot designed to work with humans in space", *Autonomous Robots*, 14(2-3), 2003, pp 179-197.

[3] A. B. Wood, T. E. Horton, and R. St. Amant, "Effective tool use in a habile agent", In *IEEE Systems and Information Engineering Design Symposium*, 2005, pp 75-81.

[4] J. J. Lockmann, "A perception-action perspective on tool use development", *Child development*, 71(1), 2000, pp 137-144.

[5] J. Piaget, *The psychology of intelligence*, Adams & Co., Totowa, Littlefield, 1969.

[6] R. W. Byrne, *The manual skills and cognition that lie behind hominid tool use*, In: A. E. Russon and D. R. Begun (eds.), The evolution of thought: Evolutionary origins of great ape intelligence, Cambridge University Press, 2004, pp 31-44.

[7] J. Carlsson, and T. Ziemke, "YAKS - Yet Another Khepera Simulator, In: Rückert, Sitte & Wittkowski (eds.), *Autonomous minirobots for research and entertainment*, HNI-Verlagsschriftenreihe, Paderborn, Germany, 2001, pp 235-241.

[8] L. A. Meeden, "An incremental approach to developing intelligent neural network controllers for robots, *IEEE Transactions on Systems, Man, and Cybernetics, Part B, Cybernetics*, 26(3), 1996, pp 474-285.

[9] N. Bergfeldt, "Towards Detached Communication for Robot Cooperation", Licenciate thesis. LiU-Tek-Lic-2005:47, Department of Computer and Information Science, University of Linköping, Sweden, 2005.

[10] G. Búason, and T. Ziemke, "Competitive Co-Evolution of Predator and Prey Sensory-Motor Systems", In: Cagnoni et al. (eds.), *Applications of Evolutionary Computing - EvoWorkshops 2003*, pp 605-615, Springer-Verlag, Berlin, 2003.

[11] G. Búason, "Competitive coevolution of sensory-motor systems", Masters Dissertation, HS-IDA-MD-02-004, Department of Computer Science, University of Skövde, Skövde, Sweden, 2002.

[12] G. Búason, N. Bergfeldt, and T. Ziemke, "Brains, Bodies and Beyond: Competitive Co-Evolution of Robot Controllers, Morphologies and Environments", *Genetic Programming and Evolvable Machines*,6(1), pp 25-51, 2005.

[13] R. St. Amant, "A preliminary discussion on tools and tool use", NCSU technical report, TR-2002-06, North Carolina State University, 2002.

[14] B. Schäfer, "Detached tool use in evolutionary robotics - Evolving tool use skills", Masters Dissertation, School of Humanities and Informatics, University of Skövde, Skövde, Sweden, 2006. Accessible at http://urn.kb.se/resolve?urn=urn:nbn:se:his:diva-12.

[15] B. B. Beck, *Animal Tool Behavior: The Use and Manufacture of Tools*, Garland Press, New York, 1980.

[16] T. Ziemke, N. Bergfeldt, G. Búason, T. Susi, and H. Svensson, "Evolving cognitive scaffolding and environment adaptation: a new research direction for evolutionary robotics", *Connection Science*, 16(4), 2004, pp 339-350.

[17] T. Susi, The Puzzle of Social Activity. The significance of tools in cognition and cooperation. Doctoral dissertation. Department of Computer and Information Science, University of Linköping, Sweden, 2006, pp 27-45.

[18] D. Kirsh, "Adapting the environment instead of oneself", *Adaptive Behavior*, 4(3-4), 1996, pp 415-452.

[19] A. S. Klyubin, D. Polani, and C. L. Nehaniv, "Tracking Information Flow through the Environment: Simple Cases of Stigmergy", In: Pollack et al. (eds.), *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems*, 4(3-4), MIT Press, 2004, pp 415-452.

[20] J. Noble and D. W. Franks, "Social Learning Mechanisms Compared in a Simple Environment", In: Standish et al. (eds.), *Artificial Life VIII: Proceedings of the Eighth International Conference on the Simulation and Synthesis of Living Systems*, 4(3-4), MIT Press, 2002, pp 379-385.

[21] H. Takeda , K. Terada, and T. Kawamura, "Artifact Intelligence: Yet another approach for intelligent robots". In *Proceedings of the 11th IEEE International workshop on Robot and Human Interactive Communication*, 2002, pp 176-182.

[22] M. Broxvall, M. Gritti, A. Saffiotti, B. S. Seo, Y. J. Cho, "PEIS Ecology: Integrating Robots into Smart Environments", *Proceedings of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Orlando, FL, 2006.