2014-01-01

Studying the Evolutionary Basis of Emotions Through Adaptive Neuroagents: Preliminary Settings and Results

Pacella, D

<http://hdl.handle.net/10026.1/10535>

10.1007/978-3-319-12745-3_5

Communications in Computer and Information Science

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Studying the Evolutionary Basis of Emotions through Adaptive Neuroagents. Preliminary settings and results

Daniela Pacella¹, Onofrio Gigliotta², Olimpia Matarazzo³, Orazio Miglino²

¹Plymouth University, Plymouth, United Kingdom
pacelladaniela@gmail.com

²University of Naples Federico II, Naples, Italy
{onofrio.gigliotta, orazio.miglino}@unina.it

³Second University of Naples, Naples, Italy
olimpia.matarazzo@unina2.it

Abstract. We propose a method to investigate the adaptive and evolutionary function of emotions and affective states – in our case, of ancestral fear - using Artificial Life and Evolutionary Robotics techniques. For this purpose, we developed a hybrid software-hardware capable to train artificial neuroagents equipped with a sensory-motor apparatus inspired on the iCub humanoid robot features. We trained populations of these agents throughout a genetic algorithm to perform a well-known neuropsychological task adapted to study emotional phenomena. The robots learnt to discriminate stressful emotional conditions (coping with “dangerous” stimuli) and no-stress conditions. Varying the network structures, the experimental conditions and comparing the outcomes we were able to delineate a very initial snapshot of behavioural and neural pre-requisite for emotional-based actions. On the other hand, we have to stress that the main contribution we brought is setting-up a methodology to support future studies on emotions in natural and artificial agents.

Keywords: Artificial Life·NeuralNetworks·Emotion·Neurorobotics

1 Introduction

The importance of studying emotions in animals for a deep comprehension of natural organisms’ behavior was denied for a long time. Behaviourists and their experimental paradigms tended to demonstrate that the desired behavior of an animal could be obtained simply through conditioning, regardless of the emotional state experienced. The fear of anthropomorphizing the research on animals also led to a neglect

of the evolutionary basis of the emotions, and their function itself was reduced to a cognitive need of labelling autonomic arousals [4].

With the rising of the neuroscientific approach, the attention shifted more to the neural aspects of behavior, and in recent years, thanks to the latest tools designed to investigate the functional organization of the cortical and subcortical structures – such as EBS and neuroimaging – the similarity of the arousal mechanisms in humans and animals became more than evident. Unfortunately, the problem of studying emotions in animals mainly lies in the difficulty of measuring and detecting variables. The methods used mostly incur into different kind of biases. In the case of the studies on the emotion of fear, for example, it is hard to isolate the correct trigger event which aroused the rat, and the measured variation of responses may be due to other factors, like daily and individual fluctuations. Also, autonomic responses are often similar among totally different kind of emotions: fear and sexual activity share the increase of cortisol level and cardiac frequency. Despite these issues, the research in the field of the neural basis of the emotions went far, and, thanks to the discover of the limbic system and of the central role of the amygdala, LeDoux in his late '90's studies came up with the first model of the conditioned fear circuit in rats [7]. The importance of subcortical structures for the genesis of fear was also proved in humans [1]. Contemporary studies also showed that decorticated rats were still capable of expressing fear [5].

Another well-known problem concerns the definition and operationalization of the emotions themselves. Regarding our terminology of fear, affective states and emotion, we refer to Panksepp's nest-layered model of the brain [11], [13, 14, 15, 16]. It states that we can distinguish three main structures inside the brain with different evolution levels and a hierarchic localization; every region is able to generate a specific set of affective states different in order and complexity from the others. Hence he distinguished three kinds of emotional circuits: *primary processes*, placed subcortically, immediate and independent from cognition (just like conditioned responses), *secondary processes*, which involve the limbic system, and *tertiary processes*, exclusively neocortical and only evidenced in humans, which include the subjective experience of complex emotions the way as we know it [13], [16]. The author, pursuing his will to investigate animal emotions, defines 7 types of proven primary processes whose circuits were found in rats: SEEKING, RAGE, LUST, FEAR, CARE, PANIC/GRIEF and PLAY. He uses capitalized letters to distinguish these basic emotional responses from the complex, cortical ones (tertiary processes) of which he has found no evidence in other species than humans [16]. These primary circuits are the proof of the existence of affective states in animals, considered as evolutionary products able to improve the performance, the adaptiveness and the survivability of an agent.

Though Panksepp's and LeDoux's models seem to rely on the totally genetic basis of the emotional responses, Panksepp himself has recently highlighted how important is to define what of an organism is genetic and what epigenetic [12]. For this purpose, he suggests the use of robotics and Artificial Life: using a computation model of the primary processes, the interactions of agents and environment can enlighten us about which role the evolution plays for the genesis of the complete experience of an emotion.

Meanwhile, studies involving artificial neural agents mainly focused on reproducing purely cognitive or motivation-driven behaviors as well as at emphasizing the complex structure of the sensory-motor apparatus, while little or no attention was given to the analysis of the models of emotional and affective states, with the exception of few preliminary works. Parisi et al., in fact, recently introduced a new kind of neuron in their networks whose activation threshold differs from the other “standard” units and able, when active, to directly communicate with the output layer and to gain priority of action, bypassing the hidden units’ computation. This immediate information processing – which leads to action - is considered of the same kind of the activation of natural organisms’ emotion circuits. As we mentioned before, in fact, these mechanisms are effective even in absence of cognition, i.e. without the cortical influence. In their work, they proved that the presence of this particular neuron, which they called “emotion unit”, significantly improved the survivability and adaptiveness of the neural agents allowing them to efficiently escape from predators or correctly prioritizing between feeding or hiding [17].

The authors of the mentioned preliminary study, however, do not provide a specification of the notion of emotion they refer to, and therefore they do not specify which emotions can be investigated with this paradigm. In our Artificial Life study we inherit the concept of “emotion unit” to set up the basis of a software able to simulate the evolution of a population of artificial agents; these robots are trained to avoid self-harming stimuli in order to spot the phylogenetic function of genetically evolved fear. We do not focus on the concept of expressing emotions, but our aim is to highlight the contribution that affective states bring to adaptation and survivability of artificial as well as natural agents in a given environment. We therefore refer to the above-mentioned concept of FEAR for setting up our experimental condition: the simulated humanoid agents (based on the iCub robot structure[8], [19]) are subjected to a cancellation task and evolved to avoid a certain type of stimulus if a specific external input is given (the presence of a light) as it becomes dangerous (causes the termination of the robot’s lifecycle). Thus, following a distinction made by Canamero [21], our work can be framed into the ‘emergent approach’ to the study of emotions.

2 Materials and methods

2.1 The neuroagent

The iCub humanoid robot is an artificial tool created to reproduce the cognitive and motor aspects of a 3-years old child. His dimensions are realistic: his height is 104 cm and weight is about 22 kg. His body has 53 d.o.f. of which 38 are placed in the upper part. His hands are extremely sophisticated and the head is provided with a visual system which allows him to follow and grasp target objects in his visual field (Fig. 1). The hardware comes up with an open-source development kit which allows to program and to transfer computer simulations on the physical robot to test it in a real environment. In order to perform the task described below, the neuroagents we used are provided with the same sensory-motor features, which grant them the ability to

explore the visual scene, to discriminate the type of items presented and to point at selected targets. To simulate our artificial robots and neurocontroller, and to set up the experimental setting, we used a modified version of the software Evorobot*, developed by Nolfi and Gigliotta and described in [10]. This tool allows the user to simulate the behavior of big populations of robots in the specified environment conditions; besides, it is possible to evolve these neuroagents with the use of genetic algorithms, selecting throughout the evolution those with the better fitness and adaptiveness. This selection process can be carried out for hundreds of descending individuals. The platform allows these simulations to be easily transferred on the real robot, placing itself as a hybrid software-hardware tool.

The sensory-motor apparatus. The neuroagent's visual system, as described above, allows the detection, identification and reaching of the targets in the experimental condition. In order to perform these actions, every artificial organism is equipped with a pan/tilt camera whose receptive area is a maximum of 350x350 pixels. The camera, moving his receptive field horizontally and vertically, allows the robot to detect the stimuli, which are distinguished according to their luminance. Also, the visual apparatus has a zoom, which represents a tool to better focus and analyse the presented stimuli on the visual scene, whose dimension is 400x400 pixels. The information perceived by the camera is computed by an artificial retina composed of 49 neurons; each neuron is able to process a visual area of 25x25 pixels.

The robot is also provided with a motor module which guides the action of selecting the targeted stimulus.

The neural controller. The network we used has a typical three-layered structure and every unit has a sigmoid activation function. The input layer is composed by 49 visual neurons and a special computational unit which allows the robot to avoid the dangerous stimuli. This unit, whose function and use have been introduced and described in [17], has an activation function different than the others and, when its threshold is reached, gets the priority over the other actions and triggers the avoiding behavior of the harming type of stimuli. This particular neuron was defined "emotional unit". The neural network hidden layer is provided with 20 hidden units, while the output layer consists of 4 different units: 2 motor neurons controlling the pan/tilt camera movements, a motor neuron controlling the magnification of the zoom and an actuator neuron which triggers the action of selecting the targeted stimulus. The basic architecture we tested is a feedforward network, though in some experiments we added recurrence at the hidden layer. Also, in some conditions we added the motor efferences to the input layer.

The adaptive algorithm. To establish the connection weights of the neural network we used a genetic algorithm that permits to evolve a population of neuroagents. In a given population of 100 agents, all robots start with random network weights and then they are repeatedly tested on the specified task; during this test set, every performance is recorded and evaluated in terms of fitness. The fitness function is a value that quantitatively estimates the behavior of an evolving robot in a given task (i.e.: exploration,

targets discrimination, predator recognizing). The values of the fitness are compared and the best 20 robots will be selected to recreate a whole new population. Each of the selected robots will generate an offspring of 5 agents with a mutation rate of the connection weights and biases of 2%. These new generation robots will constitute a new population and will be tested on the same tasks; their fitness will be compared and they will undergo the same cycle of selection. This Darwinian process is iterated for 1000 generations.

2.2 The task and the experimental setting

To induce and evaluate the conditioned fear, we built up an experimental setting in which the robot had to perform a well-known neuropsychological test, the cancellation task. The cancellation task is a widely used test and generally consists of asking the subjects to mark with a sign (e.g. using a pen or a finger) random positioned stimuli on a sheet; stimuli can be of different in shape (e.g. dots, circles, lines), colour or quantity and the indications may vary according to the specific aim of the test. Subjects can be asked to mark all the stimuli, only those in a specific position or just a part of them. This task, performed by human beings, is generally used for the diagnosis and discrimination of several kinds of spatial exploration and attentional deficits. In our experiments, the visual scene contains randomly scattered dots of two different colours: 50% are light red (high luminance stimuli) and 50% dark red (low luminance stimuli).

Fig.1. Schematic representation of the experimental setting

The number of stimuli presented changed for each trial and could vary among 2, 4, 8 or 16. When a stimulus is touched by the robot, it is cancelled from the visual scene. The neuroagents' performance is based on the number of cancelled stimuli; in half of the trials, however, light red stimuli become aversive and then caused the robot's end

of life for the specific trial. The total number of possible trials combinations is therefore 8 (4 for each condition). The change of condition is signaled by an external input (in our case, the lighting of a lamp); an example of the condition can be found in Fig.1. The artificial agents were evolved to distinguish between the two conditions and to select/cancel all stimuli in absence of the input (no-stress condition), and only dark red targets in the other case (stressful condition).

3Experiments

3.1Determining the fitness function

We carried out three main preliminary experiments in order to set up a stable environment for our platform. In the first experiment, our aim was to determine which of the used fitness functions allowed the neuroagents to reach the best performance to the task. The architecture tested in this experiment was exclusively feedforward. We compared the performance to the task of two groups of simulated robots:

Fig.2Boxplots comparing the performances of robots for each fitness function

one equipped with a fitness function which incremented with the robot's exploration ability along with his performance (number of correct targets cancelled), and the second equipped with a function which incremented only for the performance. We

evolved 10 different populations of robots for each group. In this training set, every population was evolved for 1000 generations, and every generation was tested on 20 trials, of which half in the stressful condition and the other half in the no-stress condition. For each trial, the robot could move its artificial eye for 1000 steps, at the end of which, the trial was considered finished; if a “dangerous” stimulus was cancelled during the stressful condition, the trial was terminated immediately. The 10 most evolved robots (i.e. with the highest fitness value) for each group, for a total of 20 individuals, was then subjected to a test set of 1000 trials, and their performances were compared. The results showed that the highest number of correct stimuli were cancelled by the robots of the second group: the robots who were evolved exclusively for their ability to cancel, scored higher than the others, as shown in Fig.2. In particular, in the last trial of the stressful condition the significance was $p=0,000183$ scored with the Mann-Whitney U test.

3.2 Balancing the activation of the actuator

While analyzing the activation data of the networks’ neurons, we noticed that the actuator unit (which triggered the movement of selecting and therefore cancelling stimuli) had an unexpected longer activity in terms of steps; in other words, the number of times the unit was active did not match with the number of the stimuli cancelled by the robots. This meant that the neuroagents, especially in the no-stress condition, tended to cancel target items as well as empty areas on the visual scene; this technique, in fact, allowed them to reach a higher performance. In order to punish the emergence of this sub-optimal behavior, we decided to dramatically reduce the amount of the robots’ available steps in a trial in case they cancelled an empty area. We built up two experimental conditions: in the first one, if the actuator unit was active on an empty area, the robot’s lifecycle was reduced of 250 steps, in the second one, the lifecycle was reduced of 1000 (the trial was terminated). The neural architecture used was exclusively feedforward and the fitness function incremented along with the only performance, according to the results of Experiment 1. For both conditions, we evolved for the training set 10 populations of robots for 1000 generations on 20 trials each. After the evolution, the best 10 individuals for each condition was tested on a set of 1000 trials, and results are showed in figure. There was a significant evidence that the robots within the first group (which were punished with a reduction of only 250 steps) had a better performance than the second group in the no-stress condition (respectively $p=0,000999$, $p=0,000502$, $p=0,000245$, $p=0,000187$ for the 2, 4, 8, 16 target no-stress trials). Results are showed in Fig.3.

3.3 Comparing architectures

Once all the other variables were settled, it was needed to establish which of the neural architectures best fitted with our “emotional” task. In this last experiment, we compared the performances of four different networks: I) a feedforward architecture; II)

an architecture with motor efferences; III) a recurrent network; IV) a recurrent network with motor efferences. The fitness was evaluated on the number of correct targets cancelled and if the actuator selected an empty area the trial steps were reduced of 250, in concordance with the results of the first and second experiment.

According to Panksepp's model, as described above, primary and secondary processes arise from the most ancient parts of the brain and thus without involving the cortex; these affective states stems from the brainstem and proceed through the limbic system, and do not need cognition or higher processing to be generated. Both recurrent architectures and networks provided with motor efferences have a high level

Fig.3. Boxplots comparing the performances of robots in the no-stress condition punished with a reduction of 250 or 1000 steps in case of cancellation of an empty area

of information encoding and a memory of the movements in the previous steps. These additional connections, when active, require more time for the computation of a single step and the detection of dangerous targets could be delayed; for this reason,

we expect the simplest circuit, the feedforward network, to show better reaction and to have a better performance than all the others.

We trained 10 populations of robots for each architecture for 1000 generations, each lasting 20 trials. The 10 best individuals for each population– a total of 40 individuals - were tested on 1000 trials and their performances in terms of fitness were compared. The Mann Whitney U test showed no significant difference among them ($p > 0.05$).

Fig.4. Comparison of the performances of the best robot for each architecture

We conducted further analyses extracting the single best among the best 10 individuals of each group, which we will call BestI, BestII, BestIII and BestIV for the respective architectures. We left them free to move for 1000 trials without any steps reduction and then compared their performance on correct targets, wrong targets and empty areas cancelled.

Graphics show that the better performance was reached by BestI and the difference is highly significant (Fig.4).

4 Summary and direction for future researches

We described a method to investigate the adaptive function of ancestral fear and trained artificial simulated humanoid agents to discriminate stressful and no-stress conditions in order to organize an efficient behavior and maximize their survivability. Our platform allows a rapid porting of the simulations on the physical robot iCub and therefore classifies as a hybrid software-hardware system.

Our preliminary results showed the efficiency of the simplest among the 4 tested neural architectures, confirming the importance of a rapid encoding of information for a better performance and adaptability to a given environment. Emotions and affective states provide natural agents with this literally immediate processing and reaction as soon as - like in our case - a danger is perceived, and therefore we feel like giving a contribution in favor of the theses on the phylogenetic importance of emotions. Regarding our artificial neuroagents, still much must be discovered about the true reasons lying behind their performance, and only a qualitative analysis of their behavior and movements, as well as their use of the zoom function, will reveal the difference among the results.

Other open questions regard the possibility of cutting the connections of the network with the emotion unit to watch the behavioral change of the best individuals tested. Would this result in a reversal of the obtained results? Or would the proportions be maintained?

The aim of future researches will include the replication of a standard experiment setting with various kinds of tasks which can be executed by both natural and artificial agents in order to compare their performances and to validate our network structure based on the described models of emotions.

Regarding the experimental paradigm, apart from our theoretical perspective, our tool constitutes a robotic model of the human cancellation task and its future implementation could include the reproduction of different neuropsychological phenomena.

Finally, FEAR is not the only primary process listed by Panksepp, and further simulations including different paradigms based on the other 6 primary emotions will shed light on whether the centrality of immediate processing is confirmed for all these basic affective reactions or not.

References

1. Bechara, A., Tranel, D., Damasio, H., Adolphs, R., Rockland, C., Damasio, A. R.: Double dissociation of conditioning and declarative knowledge relative to the amygdala and hippocampus in humans. *Science*, 269, 1115–1118 (1995)
2. Di Ferdinando, A., Parisi, D., Bartolomeo, P.: Modeling orienting behavior and its disorders with ecological neural networks. *Journal of Cognitive Neuroscience*, 19, 1033–1049 (2007)
3. Gigliotta, O., Bartolomeo, P., Miglino, O.: Introducing Sensory-motor Apparatus in Neuropsychological Modelization. *Artificial Intelligence and Cognition*, 80 (2013)
4. James, W.: On some Hegelisms. *Mind*, 7 (26), 186–208 (1882)
5. Kolb, B., Tees, R. C.: *The cerebral cortex of the rat*. MIT press, Cambridge (1990)
6. LeDoux, J.: The emotional brain, fear and the amygdala. *Cellular and Molecular Neurobiology*, 23, 727–738 (2003)
7. LeDoux, J.: Emotional networks and motor control: a fearful view. *Progress in Brain Research*, 107, 437–446 (1996)
8. Metta, G., Sandini, G., Vernon, D., Natale, L., Nori, F.: The iCub humanoid robot: an open platform for research in embodied cognition. *PerMIS '08 Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, 50–56 (2008)
9. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge (2001)
10. Nolfi, S., Gigliotta, O.: Evorobot: A tool for running experiments on the evolution of communication. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 297–301). Berlin: Springer (2010)
11. Panksepp, J.: *Affective Neuroscience: The Foundation of Human and Animal Emotions*. Oxford University Press, New York (1998)
12. Panksepp, J.: Simulating the Primal Affective Mentalities of the Mammalian Brain: A Fugue on the Emotional Feelings of Mental Life and Implications for AI-Robotics. In *Simulating the Mind, A Technical Neuropsychanalytical Approach*, Springer, 149–176 (2008)
13. Panksepp, J.: Affective consciousness in animals: perspectives on dimensional and primary process emotion approaches. *Proceedings of the Royal Society*, 277, 1696 2905–2907 (2010)
14. Panksepp, J.: Cross-Species Neuroaffective Parsing of Primal Emotional Desires and Aversions in Mammals. *Emotion Review*, 5(3), 235–240 (2013)
15. Panksepp, J., Fuchs, T., Iacobucci, P.: The basic neuroscience of emotional experiences in mammals: The case of subcortical FEAR circuitry and implications for clinical anxiety. *Applied Animal Behaviour Science*, 129(1), 1–17 (2011)

16. Panksepp, J., Biven, L.: *The Archaeology of Mind: Neuroevolutionary Origins of Human Emotions*. W. W. Norton & Company, New York (2012)
17. Parisi, D., Petrosino, G.: Robots that have emotions. *Adaptive Behavior*, 18(6), 453-469 (2010)
18. Ponticorvo, M., Miglino, O.: Encoding geometric and non-geometric information: a study with evolved agents. *Animal Cognition*, 13(1), 157-174 (2010)
19. Schmitz, A., Pattacini, U., Nori, F., Natale, L., Metta, G., Sandini, G.: Design, Realization and Sensorization of the Dexterous iCub Hand. *International Conference on Humanoid Robots*, 186-191 (2010)
20. Vingiano, W.: Pseudoneglect on a cancellation task. *International Journal of Neuroscience*, 58(1-2), 63-67 (1991)
21. Canamero, L.: Emotions understanding from the perspective of autonomous robots research. *Neural Networks*, 18, 445-455 (2005)