

The evolution of pain

Alberto Acerbi¹, Domenico Parisi¹

¹ Institute of Cognitive Sciences and Technologies, National Research Council
44, Via San Martino della Battaglia, 00185, Rome, Italy
{alberto.acerbi, domenico.parsi}@istc.cnr.it

Abstract. We describe two simple simulations in which artificial organisms evolve an ability to respond to inputs from within their own body and these inputs themselves can evolve. In the first simulation the organisms develop an ability to respond to a pain signal caused by body damage by stopping looking for food when they feel pain since resting while the body is damaged accelerates healing of the body and increases the individual's survival chances. In the second simulation the pain signal itself evolves, that is, the body develops a tendency to send pain signals to the nervous system when the body is damaged. The results are discussed in terms of an internal robotics in which the robot's body has an internal structure and not only an external morphology and the neural network that controls the robot's behavior responds to inputs both from the external environment and from within the body.

Keywords: artificial life, neural nets, internal robotics, pain.

1 Introduction

The nervous system responds both to information from outside the body (external environment) and to information from the organs and structures that lie inside the body (internal environment). In both cases the nervous system tends to respond to the incoming information in ways that increase the survival and reproductive chances of the organism. However, there is an important difference between the external environment and the internal environment. The external environment has properties that exist independently of the organism. It is true that the organism's sensory organs filter the environment in ways that depend on the organism and that, by moving its body or body parts, the organism can determine what portions of the environment will send information to its sensory organs at any given time. However, the intrinsic properties of the environment and the fact that these properties have an impact on the organism's nervous system are something for which the organism has no responsibility. The internal environment is different. In the case of the internal environment what must evolve is not only an ability of the organism's nervous system to respond in appropriate ways to the information originating from the internal structures of the body but also the fact itself that these internal structures can send specific kinds of information to the nervous system. In other words, in the interactions between the nervous system and the rest of the body, everything must evolve.

In most simulations using neural networks and genetic algorithms what evolves is the manner in which organisms respond to the input but the input itself does not evolve. The input has properties that are the physical consequence of the intrinsic nature of the external environment. But consider physical pain. Physical pain is input from within the organism's body caused by some damage, or potential damage, to the body. The organism's nervous system must respond to pain in ways that tend to increase the organism's survival and reproductive chances. A simple adaptive response to physical pain, often found in animals, is to stop moving so that whatever the damage to the body that causes pain it can more easily heal spontaneously. In the case of pain, not only the nervous system's response to pain must evolve (to stop moving) but pain itself as a signal to the nervous system must evolve. The body must evolve a tendency to translate physical damage that occurs in some of its parts into pain. In fact, people who are unable to feel pain tend to die before they reach adulthood. Therefore, pain as an input to the nervous system can be considered as an adaptation [8, 15, 16].

In this paper we describe some simple simulations of the evolution of pain. In the first simulation a population of artificial organisms evolve an ability to find food in the external environment but at the same time they stop moving and looking for food when they feel pain since resting while the body is damaged accelerates healing of the body and increases the individual's survival chances. In a second simulation pain itself evolves, that is, the body evolves a tendency to send pain signals to the nervous system when the body is damaged. In the first simulation inherited genotypes encode only the properties of the organisms' nervous systems, i.e., the connection weights that allow the neural networks to respond appropriately to incoming signals from either the external or the internal environment. In the second simulation genotypes also encode some properties of the organisms' body beyond the nervous system. They encode a tendency of the body to send pain signals to the neural network to which the neural network can respond appropriately.

2 Simulations

2.1 Evolving the ability to respond to pain

A population of 100 organisms lives in an environment containing food elements. Each individual lives in its own copy of the environment for 10 epochs of 50 time steps each (input/output cycles of the neural network controlling the organism's behavior). The environment is a continuous square with 5 randomly located food elements. At the beginning of each epoch the individual is positioned in a randomly selected location with a randomly selected orientation. For independent reasons and at unpredictable intervals each individual has a 0.01% probability at each time step to incur in some physical damage to its body. The damage lasts for 20 consecutive time steps and it then recovers spontaneously. The nervous system of the organism is simulated by a three-layered neural network with three units in the input layer and

two units each in the internal and output layers. Two input units encode the distance of the nearest food element from the organism and the angle between the food element and the individual's current orientation. (The organism has a perceptual field of 360 degrees.) These two units are connected to the two internal units which in turn are connected to the two output units. The third input unit encodes the presence (1) or absence (0) of pain which is felt when the organism's body is physically damaged. This input unit is directly connected to the two output units. The two output units encode the organisms' movements in the environment. One output unit, the orientation unit, encodes the angle with which the individual's orientation changes with respect to its present value, and the other unit, the displacement unit, encodes the distance covered by the individual in its present orientation.

At the beginning of the simulation the neural networks of all the organisms have connection weights with a value randomly selected in the interval between -5 and $+5$. If an organism reaches a distance which is smaller than 0.01 from a food element, the organism eats the food element and its fitness is increased by one unit, while the food element which has been eaten is replaced by a new one in a randomly selected location of the environment. If an organism displaces itself in the environment when the pain unit is 1, that is, its second output unit has a value higher than 0.01 , the organism's fitness is decreased by 2 units. At the end of life the 20 individuals with the highest fitness are selected for reproduction, and they generate 5 offspring each which inherit their connection weights, with a probability of 0.01 of mutating each connection weight by adding or subtracting a value randomly selected within the same range between -5 and $+5$. The $20 \times 5 = 100$ new individuals constitute the second generation. The simulation is terminated after 1000 generations. All the results that are cited in the paper are the average results of 10 replications of the different simulations.

Figure 1 show how the fitness of the average individual and the fitness of the best individual increase across the 1000 generations. Figure 1 also shows the curves for the average and best individuals of a control simulation in which there is no physical damage and therefore no pain signal.

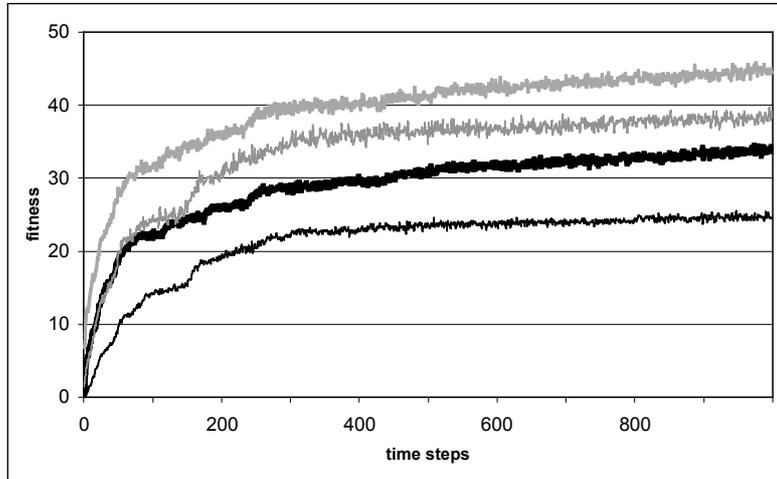


Fig. 1. Average (tick black line) and best (tick gray line) fitness for Simulation 1 in which the organisms evolve an ability to both look for food when they are in a healthy condition and do not feel any pain and to stop moving and ignore food when their body has some physical damage and they feel pain. Average (thin black line) and best (thin gray line) fitness for a control simulation in which there is no physical damage and therefore no pain signal.

Figure 1 shows that the organisms evolve an ability to look for food both when their body can be damaged and they have to stop looking for food and when body damage cannot occur. Inevitably, the total quantity of food they are able to eat is lower in the first than in the second condition. However, the organisms are very good at stopping moving only when they feel pain. This is shown in Figure 2. The organisms stop moving in 90% of the cycles in which they receive the pain stimulus but they stop moving only in 15% of total cycles, that is, approximately, the percentage of cycles in which they in fact receive pain stimuli.

The inhibitory function of the pain signal is realized by evolving strong inhibitory (negative) connection weights for the connection linking the pain input unit to the output unit which controls speed of movement. If we examine how the weights of the two connections linking the pain input unit to the two output units change across generations, we find that while the weight value of the connection to the first output unit (angle of turning) oscillates around a value of zero, the weight value of the connection to the second output unit (speed of movement) becomes strongly negative (inhibitory) after a few generations (Figure 3). Since when the body is damaged, it is displacing oneself in the environment, not turning, which is costly in terms of fitness, it is reasonable that only the second connection weight becomes strongly negative, thereby inhibiting the individual's displacements in the environment when the body is damaged and the pain signal is on, that is, the pain input unit has a value of 1.

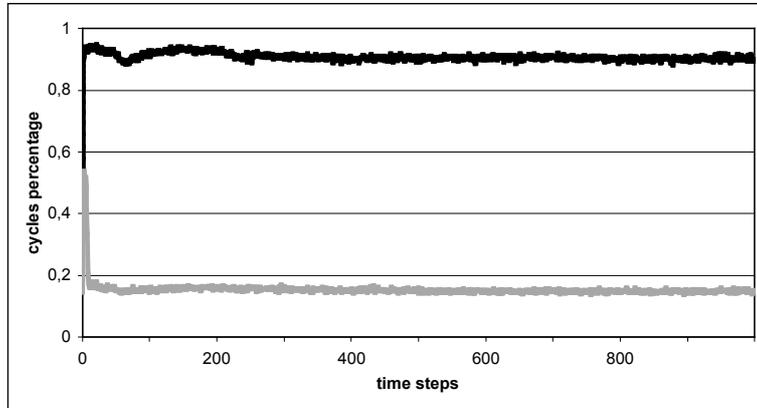


Fig. 2. Percentage of input/output cycles in which the organisms stop moving when they feel pain (black line) and total percentage of cycles in which they stop moving (gray line).

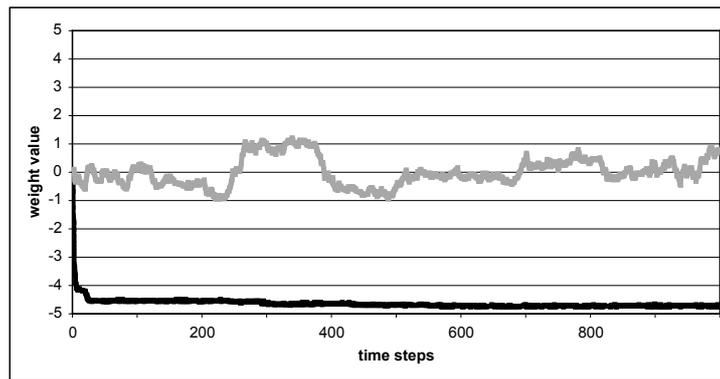


Fig. 3. Weight value of the connections linking the pain input unit to the first output unit (angle of turning - gray line) and to the second output unit (speed of movement - black line) across 1000 generations.

2.1 Evolving the ability to respond to pain

Simulation 1 has shown that, given an input signal that co-varies with physical damage to the body, the organisms learn to stop moving when this signal from the body is on because moving when there is physical damage to the body decreases their reproductive chances. But, as discussed in the Introduction, inputs from inside the body, unlike inputs from the external environment, are not automatic consequences of

the independent physical structure of the internal environment but they must co-evolve together with the body. The body itself must “learn” to send signals to the nervous system and must decide which specific signals to send and in correspondence with which specific states of the body. This is what happens in Simulation 2.

In Simulation 1 the decision to activate the pain input unit when there is physical damage to the organism’s body but not otherwise was hardwired. In the new simulation the pain signal from the body can be seriously disturbed by random noise. We add a new “gene” to each individual’s genotype (that encodes the connection weights of the individual’s neural network) and this new gene can have a value ranging from 0 to 1. The value of the gene determines the amount of random noise which is added to the pain signal. If the noise gene has a value of 0, the pain signal perfectly informs the nervous system about the state of the body. On the other hand, with increasing values of the noise gene, the pain signal becomes progressively less informative because progressively more noise is added to the signal. With high values of the noise gene, the nervous system does not know if the body is damaged or not.

At the beginning of the simulation there is no useful pain signal from the body to the nervous system. We simulate this by assigning very high values to the noise gene for the pain signal, between 0.9 and 1.0. This implies that fitness cannot be very high since the organisms simply do not know when it is appropriate to move and to look for food and when it is appropriate not to move and to rest. However, the individuals that reproduce transmit their entire genotype to their offspring, which includes both the neural network’s connection weights and the noise gene for the pain signal. Like the connection weights, the inherited value of the noise gene is randomly mutated by replacing with a probability of 0.01 the gene’s current value with a new value randomly selected between 0 and 1.

The results show that after about 200 generations the gene’s average value tends to be less than 0.1 (Figure 4). This means that the organisms’ nervous system can now know with sufficient precision when the body is physically damaged and when it is not, and it can respond to the first information by stopping moving and ignoring food and to the second information by moving and eating.

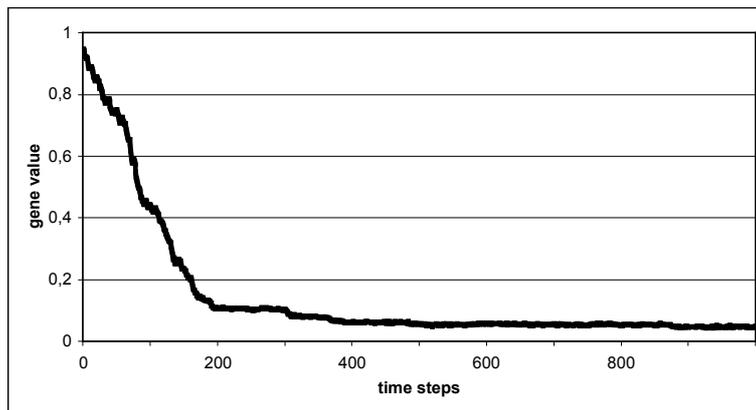


Fig. 4. Average value of the noise gene for the pain signal in Simulation 2.

The fitness value at the end of the simulation is similar to the fitness value at the end of Simulation 1 (Figure 5), showing that an effective integrated system has emerged: the organisms' body has evolved an appropriate signal that tells the nervous system when the body is damaged and the nervous system has evolved an appropriate response to this signal (resting).

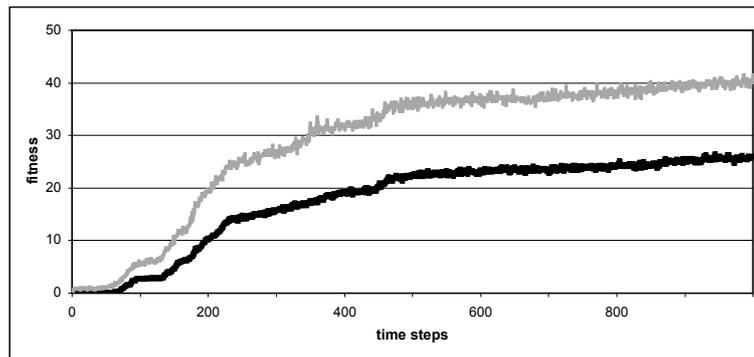


Fig. 5. Average (black line) and best (gray line) fitness for Simulation 2.

As in Simulation 1, the organisms tend to stop moving when they feel pain and to move and search for food when there is no pain signal. The appropriate performance level is reached somewhat more slowly than in Simulation 1 because evolution has to solve two problems rather than only one as in Simulation 1.

3 Discussion

Pain clearly does have adaptive value [8]. Clinical cases demonstrate that people who don't experience pain die at an early age, from accidental injuries or, more frequently, as a consequence of irreversible damage to their joints due to the fact that, since they don't feel pain, they do not change frequently enough the posture of their joints. One clinical condition is syringomyelia, a gradual deterioration of some portions of the spinal cord that selectively eliminates pain in various part of the body, especially the hands [6].

We have been able to evolve simple artificial organisms that not only have nervous systems that respond to inputs from the environment with the appropriate behaviors but also have bodies that send the appropriate inputs to the nervous system. When the body incurs physical damage, it is appropriate from the point of view of the individual's reproductive chances to stop moving even if moving is necessary to find food and the individual must move and look for food when its body is healthy. But for this more complex behavior to emerge the organism's nervous system needs to know the current state of the organism's body. We have shown that not only can our artificial organisms evolve an ability to modulate their behavior so that they move and

look for food when there is no signal that their body is damaged and stop moving and rest when there is a pain signal indicating that the body is damaged, but also that their body can evolve an ability to signal appropriately to the nervous system what is its current state. (For a description of the actual neurobiology of pain, i.e., of an organism's nociceptive system, see

The body provides the nervous system with a great variety of important signals. For example, hunger and thirst are signals sent by the body to the nervous system and informing the nervous system about the current quantity of energy and liquids present in the body, respectively. On the basis of these signals the nervous system can take the appropriate action by producing either behaviors that reintegrate energy (eating) or behaviors that replenish liquids (drinking). Artificial organisms living in an environment that contains both food and water must know what is the level of both energy and liquids currently existing in the body in order to decide at any given time whether to look for food and ignore water or to look for water and ignore food. In fact, it has been shown that artificial organisms not only can evolve nervous systems that respond appropriately to hunger and thirst signals from their body [4, 13] but can also evolve bodies that send the appropriate hunger and thirst signals to the nervous system [3].

The simulations that address how artificial organisms evolve an ability to respond to inputs from inside the body and how these inputs themselves can evolve, can be considered as part of an internal robotics [14] that needs to be developed in addition to the more traditional external robotics if we want to understand more completely the behavior of organisms. Another example of internal robotics are artificial organisms that respond to light in the environment by moving in search of food and to darkness by stopping moving (sleeping). However, if these organisms enter a dark cave, they may stop moving for ever. The problem is solved by having these organisms evolve an internal "biological clock" that tells their nervous system when it is daytime and when it is nighttime independently of sensory input from the external environment [9]. (For other explorations of internal robotics, see [5])

The simulations described in the present paper have addressed physical pain associated with physical damage to the body and have shown that physical pain can be interpreted as an adaptive signal which is sent by the body to the nervous system and which evolves as part of the evolution of the entire body. An interesting question is whether the same interpretation and the same approach that we have used to simulate the evolution of physical pain can be applied to psychological pain. Physical pain signals a physical damage to the body, and it is the rest of the body, outside the nervous system, which is at the origin of the pain signal and sends the pain signal to the nervous system. (Physical damage to the nervous system, at least to the central nervous system, generally does not cause pain.) In contrast, signals of psychological pain (grief) seem to be self-generated within the nervous system. But, aside from the physical origin of the different pain signals, the more interesting question is what is the adaptive significance of psychological pain signals and what is the adaptive response to these signals. In the simple scenario of our simulations the appropriate response to physical pain is resting. When they feel physical pain our organisms simply stop moving since resting facilitates the spontaneous recovery of the body from physical damage. In fact, resting appears to be a simple but very fundamental adaptive response to physical pain which exists in many non-human animals,

although of course more sophisticated and differentiated adaptive reactions to physical pain exist even in non-human animals and certainly in human beings. Examples include learning to avoid whatever caused the pain, acting on the specific part of the body which is in pain (which of course requires a spatially more rich information from the body), and going to a doctor.

But what is the adaptive response to psychological pain which justifies the evolutionary emergence of signals of psychological pain self-generated within the nervous system? Signals of psychological pain may cause in the individual that feels the pain behaviors that allow the individual to escape from situations that have negative consequences for the individual's survival chances or to avoid future situations with these consequences. This may apply to the psychological pain felt by young helpless individuals when they are separated from their mother [2] but it can generalize to other cases of psychological pain in adults. (For a general evolutionary account of various forms of psychological pain, cf. [10, 11, 12]. For a different position that views psychological pain as a by-product of other processes, see [1].)

Another interesting aspect of pain that could be addressed with evolutionary simulations is the anticipatory value of pain. Pain is not only a signal that allows the organism to escape from current conditions that are negative from an adaptive point of view (e.g., moving when the body is damaged or being separated from parent) but it is by itself a condition that the organism tries to escape from or to avoid. Hence, pain appears to have an anticipatory value. By trying to escape from or avoid pain, the organism escapes from or avoids future conditions that may have negative consequences from an adaptive point of view.

The anticipatory value of pain applies especially to psychological pain. While physical pain occurs only if there is actual damage to the body, psychological pain is something that can be learned. Initially neutral stimuli can become psychologically painful with learning. (Of course physical pain itself can cause psychological pain. For the interactions between physical and psychological pain, see [7]). This allows psychological pain to assume a great importance from an anticipatory point of view, alerting the organism to future negative conditions and allowing the organism to act preventively in order to avoid those conditions. By realizing simulations in which pain does not only evolve but can also be learned during life, one could address this important property of psychological pain.

A final problem posed by pain is its "felt" nature. Pain, both physical and psychological, is not just a process or an event which takes place in our body/nervous system but is something which is subjectively "felt". Can our evolutionary simulations address the problem of the "felt" nature of pain? This a complex question that cannot really be addressed here. However, we can indicate a direction that our simulations could take to address this question.

At any given time many different inputs arrive to an organism's nervous system and the organism cannot simultaneously respond to all of them. It has to choose. The "felt" nature of pain can be interpreted as a mechanism for guaranteeing that some particular input is given prominence in determining the organism's behaviour. In general terms, "felt" states are states that "speak louder" in order to be heard (responded to) by the organism. From this point of view pain can reveal its adaptive nature if we assume that the behaviour which is generated in response to the particular input to which pain gives prominence is, in the circumstances, the most important one

from an adaptive point of view. The mechanism through which pain plays this role of giving prominence can basically be the inhibitory mechanism that has emerged in our Simulation 1. Pain stimuli are inputs that inhibit other inputs (in our simulations, the input from food) which would produce other behaviours (in our simulations, the behaviour of moving and searching for food) and in this way they make it possible for the organism to produce a behaviour which, in the circumstances, is more important from an adaptive point of view (in our simulations, the behaviour of resting when the body is damaged).

References

1. Archer, J.: *The Nature of Grief. The Evolution and Psychology of Reaction to Loss*. Routledge, London (1999)
2. Besson, J. M.: The neurobiology of pain. *The Lancet* 353 (1999) 1610-1615
2. Bowlby, J.: *Attachment and Loss. Volume 2. Separation: Anxiety and Anger*. Hogart, London (1975)
3. Cangelosi, A., Parisi, D., Nolfi, S.: Cell division and migration in a 'genotype' for neural networks. *Network* 5 (1994) 497–515
4. Cecconi, F., Parisi, D.: Neural networks with motivational units. In: Meyer, A.J., Roitblat, H.L., Wilson, S.W. (eds.): *From animals to animats 2: Proceedings of the 2nd International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge (1993) 346–355
5. Fellous, J.M., Arbib, M. (eds.): *Who needs Emotions? The Brain Meets the Robot*. Oxford University Press, Oxford (2005)
6. Hendry, S.: Pain. In: Wilson, R., Keil, F. (eds.): *MIT Encyclopedia of Cognitive Sciences*, MIT Press, Cambridge (1999) 622–624
7. Loeser, J.D., Melzack, R.: Pain: an overview. *The Lancet* 353 (1999) 1607–1609
8. Melzack, R.: *The Puzzle of Pain*. Basic Books, new York (1973)
9. Mirolli, M., Parisi, D.: Artificial organisms that sleep. In: Banzhaf, W., Christaller, T., Dittrich, P., Kim, J.T., Ziegler, J. (eds.): *Advances in Artificial Life. 7th ECAL Conference*, Springer, New York (2003) 377–386
10. Nesse, R.M.: What good is feeling bad? The evolutionary benefits of psychic pain. *The Sciences nov./dic.* (1991) 30–37
11. Nesse, R.M.: The Evolution of Hope and Despair. *Social Research* 66 (1999) 429–469
12. Nesse, R.M.: Is depression an adaptation? *Archives of General Psychiatry* 57 (2000) 14–20
13. Parisi, D.: Motivation in artificial organisms. In: Tascini, G., Esposito, V., Roberto, V., Zingaretti, P. (eds.): *Machine Learning and Perception*, World Scientific, Singapore (1996) 3–19
14. Parisi, D.: Internal robotics. *Connection Science* 16 (2004) 325–338
15. Wall, P.: *Pain. The science of suffering*. Weidenfeld and Nicolson, London (1999)
16. Willis, W. D.: *The Pain System: The Neural Basis of Nociceptive Transmission in the Mammalian Nervous System*. Kragel, Basel (1985)