Robot pain: a speculative review of its functions

Carme Torras
Institut de Robòtica I Informàtica Industrial (CSIC-UPC)
http://www.iri.upc.edu/people/torras

"There are no gains, without pains..."
Benjamin Franklin, The Way to Wealth (1758)

"Their minds are so narrow that can't hold the slightest doubt"
Emil Cioran

1. Introduction

Nowadays robots exhibit a wide variety of cognitive abilities, among them that of learning from bad experiences. Does this mean that robots may have doubts about what course of action to take and ultimately be able to experience pain?

Before trying to answer this question, let me quote the Nobel prize Herbert A. Simon, who in the late sixties claimed that we should not push biological inspiration too far when designing artificial creatures, since the best natural mechanisms may not be the best artificial ones [20]. Wheels, airplanes' non-flapping wings, and calculators have often been mentioned as examples of artificial solutions considerably different from their natural equivalents, and more performant according to certain criteria. The resources available to engineering design depart a lot from those in nature, and not just when it comes to materials, but also in the number of instances and spendable time. Robots have been incrementally designed along sixty years of technological research, instead of being the outcome of thousands of years of evolution.

Thus, in case of existence, *robot pain* and *robot cognition* must be of a different nature than human ones, and if these labels are used, it is because the underlying functions play a similar role as those of their human counterparts.

With this note of caution in mind and accepting that biological plausibility adds no special value from an engineering viewpoint, there is no doubt that human biology constitutes a great source of inspiration for robotics researchers, since it provides existence proofs of many useful physical and cognitive mechanisms. Thus, robots often have an anthropomorphic anatomy with legs, arms and hands having as many degrees of freedom as the human ones, their actuators mimicking the flexibility and power of biological muscles, with visual, auditive, tactile as well as proprioceptive sensors, and with a hierarchical control system whose uppest layer is a symbolic decision maker that governs lower layers often in the form of neural networks encoding sensorimotor mappings.

In trying to unravel the roles pain may play at all these levels, we adopt the distinctions between types of pain made by de Vignemont [4]. A first distinction is between chronic pains, occurrent pains, and pain expertise based on prior experiences of pain. Since chronic pains do not have adaptive value and thus cannot fulfill a function in robot development, we will not consider them further. *Occurrent pains* can be used as punishment within reinforcement learning schemes, as described in Section 2, whereas *pain expertise* can be incorporated in a motivational/cognitive system modulating reactions to stimuli, as elaborated in Section 3. One can further distinguish between merely unpleasant and excruciating pains, upon which the motivational modulation will act differently.

Another useful distinction is between *exogenous* and *endogenous* pains. Although it is clear that the underlying criterion is whether the cause is external or internal to the organism, in order to provide concrete examples of each we need to enter the slippery terrain of defining what we mean by robot pain. A short digression may help. This is an excerpt of a fictional dialog [5] in which "Russell" argues for the importance of a clear definition of pain to make scientific progress, while "Edison" takes the pragmatic view of the inventor who just wants to build robots whose pain serves a function:

RUSSELL: How do we define self and pain in ways that even begin to be meaningful for a machine? For example, a machine may overheat and have a sensor that measures temperature as part of a feedback loop to reduce overheating, but a high temperature reading has nothing to do with pain.

We agree: this would correspond to nociception, not pain, at least in humans.

EDISON: I disagree! Overheating is not human pain for sure but certainly "machine" pain! I see no problem in defining self and pain for a robot. The self could be (at least in part) machine integrity with all functions operational within nominal parameters. And pain occurs with input from sensors that are tuned to detect non-nominal parameter changes (excessive force exerted by the weight at the end of a robot arm).

Paradoxically, we agree with this view too. A percept that compromises the normal functioning of the robot could be interpreted as robot pain.

RUSSELL: [..] my earlier examples were to make clear that "pain" and detection of parameter changes are quite different. If I have a perfect local anesthetic but smell my skin burning, then I feel no pain but have sensed a crucial parameter change. True, we cannot expect all aspects of human pain to be useful for the analysis of robots, but it does no good to throw away crucial distinctions we have learned from the studies of humans or other animals.

Of course, we intend to keep crucial distinctions, and we note that the word "crucial" has a self-referential meaning in the paragraph above, as it provides the key to our pragmatic definition of robot pain: a crucial parameter change that diminishes robot competence and degrades its performance.

EDISON: [..] Inspired by what was learned with fear in rats, a roboticist would say "OK! My walking robot has analogous problems: encountering a predator—for a mobile robot, a car or truck in the street—and reacting to a low battery state, which signals the robot to prepare itself for functioning in a different mode, where energy needs to be saved." Those two robot behaviors are very similar to the rat behaviors in the operational sense that they serve the same kind of purpose. I think we might just as well call them "fear" and "pain." I would argue that it does not matter what I call them—the roboticist can still be inspired by their neural implementations and design the robotic system accordingly.

We will follow this biologically-inspired, operational approach here and, adhering to our above-mentioned pragmatic definition of robot pain, colliding with an obstacle or getting too close to a light source so that the camera saturates would be examples of exogenous pain, while an endogenous pain would involve proprioceptive sensors detecting, for instance, that a robot arm is at a singular configuration, thus lacking one degree of freedom, or simply that a joint has reached its limit, thus losing manipulability. All these situations have in common that robot competence decreases and performance degrades.

It is worth mentioning that endogenous pains can contribute to the acquisition of a body schema by the robot (a kinematic model of its mobility), which is loosely reminiscent of the development of a sense of bodily ownership in humans, as described in Section 2.2. Although some authors relate the discovery of the body boundaries to the distinction between self and others, and ultimately to the emergence of a conscious self, we share the view of Hoffmann et al. (2010) that robots have not yet reached the level of competence where notions like conscious representations can be investigated in a grounded fashion.

Going back to the initial question, in this chapter we will try to unravel the role pain (and choice) may play in developing robot competences. We can pinpoint at least three adaptive functions where pain plays a role in the biological world and which have been incorporated into robots, namely learning through reward/punishment (both at the sensorimotor and cognitive levels), providing intrinsic motivation, and engaging in empathic interaction. While in punishment and aversive motivation the robot can be said to "experience" pain, although in its own artificial way, the interaction context entails perceiving pain in humans and (up to now) faking empathy so as to give them comfort, thus being of a very different nature. In what follows we will concentrate on the former contexts in which robots "experience" pain, which can all be formulated in the common framework of reinforcement learning.

2. Occurrent pain as punishment in reinforcement learning

In the biological world, occurrent pain clearly serves an adaptive purpose by eliciting escape responses and thus protecting organisms from further danger. Taking a design perspective, adaptivity could be dispensable for robots that perform repetitive tasks in predefined environments (such as production lines), but it is a *sine qua non* if they are to work and survive in human, dynamic scenarios, where there is no expert programmer at hand. Among the repertoire of artificial adaptive mechanisms ranging from unsupervised to supervised ones [24], reinforcement learning (RL) algorithms are the most appropriate in the aforementioned scenarios, where evaluating the performance of a behavior is possible but the optimal behavior is unknown.

Reinforcement Learning (RL) refers to the process of improving performance through trial and error. The simplest RL algorithms rely on the intuitive idea that if an action is followed by an improvement in the robot's state, then the tendency to produce such action should be strengthened in that situation, whereas it should be weakened if a negative effect like pain occurs. This is the classic mechanism of reward/punishment. More formally, RL algorithms entail progressively building a mapping from situations to actions so as to maximize a scalar reward or reinforcement signal. The robot is not told which action to take, as in supervised learning, but instead must discover which actions yield the highest reward by trying them. In more complex cases, actions may affect not only the immediate reward, but also the next situation, and through that all subsequent rewards. These two characteristics—trial-and-error search and delayed reward—are the two most important distinguishing features of reinforcement learning [22].

Figure 1(a) shows the classic diagram of an RL agent (e.g., a robot controller) interacting with its environment. The agent generates actions in the context of sensed states of this environment, and its actions influence how the environment's states change over time. The environment contains a "critic" component that, at each time step, provides the agent with a numeric evaluation of its ongoing behavior. The term critic is used, instead of "teacher", because in supervised learning algorithms a teacher provides more informative instructional information, such as directly telling the agent what its actions *should have been* instead of merely scoring them.

In the biological world, the critic's evaluation corresponds to what behavioral psychologists call *primary reward*, namely that encouraging behavior directly related to survival and reproductive success, such as eating, drinking, and escaping. The mapping from states to rewards is called a *reward function*, which is an essential component of RL algorithms since it implicitly encodes the problem the agent must learn to solve. Pain takes the form of negative reward.

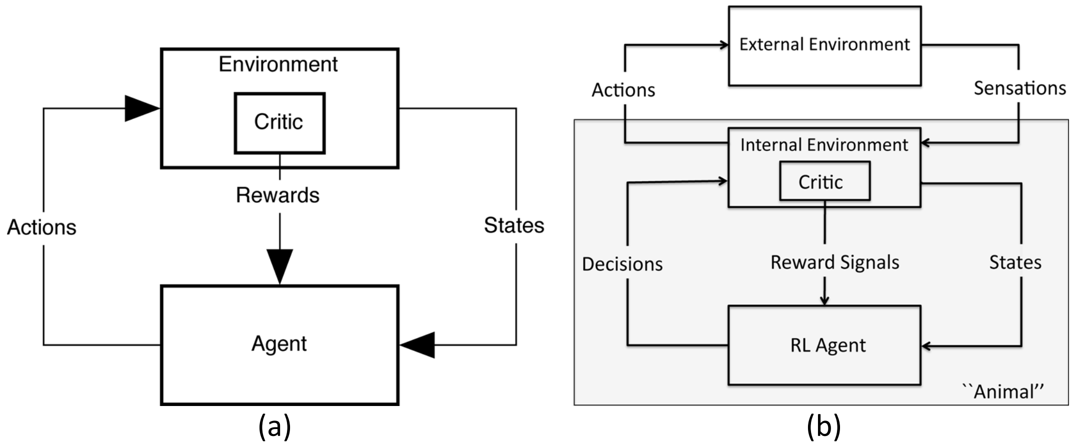

Figure 1. *Agent-environment interaction in reinforcement learning [2]. (a) Primary reward signals are supplied to the agent from a "critic" in its environment. (b) A refinement in which the environment is divided into an internal and an external environment, with all reward signals coming from the former. The shaded box corresponds to what we would think of as the animal or robot.*

As mentioned above, the objective of a complex agent is to act at each time step so as to maximize not the immediate reward, but the total reward it expects to receive over the future (called the *expected return*), which is usually a weighted sum in which later rewards are weighted less than earlier ones. Because the agent's actions influence the environment's state over time, maximizing expected return requires the agent to control the evolution of its environment. This is very challenging, since the agent might have to sacrifice short-term reward in order to achieve more reward over the long term. Some pains are bearable if the expected return is high enough.

The simplest RL agents attempt to achieve this objective by adjusting a *policy*, which is a rule that associates actions to observed environment states. A policy corresponds to a stimulus-response (S-R) rule of animal learning theory. But RL is not restricted to simple S-R agents: more complex RL agents learn models of their environments that they can use to make plans about how to act appropriately. This entails the capacity to predict and choose between alternative courses of action.

Analogously, RL in Robotics is applied in two rather different forms depending on whether it takes place at the sensorimotor or cognitive levels. Sensorimotor adaptation is implemented by building mappings from stimuli to proper movements, while cognitive learning entails constructing symbolic representations to guide decision-making. In the following two subsections we will describe how associative S-R agents have been used to learn robot sensorimotor mappings and a body schema, respectively, whereas symbolic RL agents for learning to plan will be described in Section 3.2. A detailed survey of the challenges and successes in the application of RL to Robotics is provided in [9].

2.1. Reinforcement learning of sensorimotor mappings

Motion control, both in biological and technological systems, relies strongly on sensorimotor mappings. These mappings vary depending not only on the nature of the involved sensors and actuators, but also on the goal pursued. Hand-eye coordination and stable walking are skills that may be acquired through reinforcement learning, their underlying sensorimotor mappings involving visual and proprioceptive signals, respectively, and their reward functions relying on somewhat delayed reinforcement in the form of success (object grasped) or pain (falling down).

Despite their apparent disparity, these mappings have in common an underlying highly nonlinear relation between a continuous (often hard to interpret) input domain and a continuous motor domain; a relation that in most cases would be impossible to derive analytically. Thus, such relation needs to be approximated through associative learning. Neural network architectures have the versatility to encode a large range of nonlinear mappings and, through RL, they have proven adequate to handle the massively-parallel task of relating perception patterns to motor commands.

For further details, the reader is referred to a review of neural learning algorithms used to approximate sensorimotor mappings for the control of articulated robots [25]. Similar procedures apply to wheeled robot navigation [14], where a mapping from proximity sensor readings to obstacle avoidance motions is learnt by using as reward function the pain (or punishment) derived from collisions with environmental obstacles.

2.2. Acquisition of a body schema

In this volume, de Vignemont [4] relates pain to the learning of the body boundaries and ultimately to the development of a bodily self. Along the same line but without referring to pain, Haselager et al. [6] highlight the importance of sensing one's own movements for the development of a nonconceptual sense of self. These authors claim that proprioception and kinesthesis are essential in this development, and make a plea for robots to be equipped with a richer sense of proprioception, so as to advance in the understanding of creatures acting in the world with a sense of themselves.

In neurosciences, the correlation of proprioceptive sensory information, e.g. joint configurations, with the visible shape of the body has been termed "body schema" and is assumed to provide an unconscious awareness of the current body state. Experiments with both macaque monkeys and humans have shown that the body schema is not congenital but learnable [11], and may not be a localized pattern but rather be distributed across several groups of connected neurons that represent opportunistically-learned manifolds spanning several brain regions [21].

Hoffmann et al. [7] provide a detailed review of robotics works tackling the learning of a body schema, which they define to be a sensorimotor representation of the body used for action. As an example, for arm robots such representation is usually restricted to an inverse kinematics/dynamics mapping relating the gripper's pose (position and orientation) and velocity to the arm's joint angles and torques. Then, according to this definition, the body schema can be learned in the same way as the sensorimotor mappings in the previous subsection by using pose/velocity as input and the robot kinematic parameters as output.

Besides this parameter adjustment of phylogenetic skills through learning, biological creatures show an impressive capacity to adapt their body scheme to newly acquired abilities (e.g., tool usage) as well as to abrupt changes in their bodies (e.g., limb loss). For instance, a monkey's proximal visual receptive field was enlarged by the length of a tool after being trained to use it [11]. This led to hypothesize that the tool itself became incorporated into the monkey's own body schema [3]. Similar conclusions were drawn from experiments with human patients who suffered from phantom pain after having lost a limb.

Not surprisingly, replicating this capacity has arisen a lot of interest in the robotics community, since making robots adaptive to changes in their own geometry due to damage, wear-and-tear or tool replacement increases their robustness and widens their range of application. Hence, neural learning algorithms have been used to learn not the whole kinematics mapping from scratch, as previously done, but only the deviations from the nominal kinematics embedded in the original robot controller [18], resulting in a significant speed-up that permits online adaptation of the body schema, a crucial feature for autonomous robots that need to operate for a long time without assistance (e.g., in space). Likewise, the body schema of humanoid robots has been extended to encompass a tool through training [26] and RL algorithms relying on a self-model have been implemented on an hexapod robot that needed to adapt to leg removal, broken legs and motor failures [10].

As a final note, let us mention that hand in hand with the development of the body representation, infants acquire a notion of body ownership and agency, which at the present time goes beyond the scope of the functional approach to robot pain adopted in this chapter.

3. Pain expertise as a form of emotion and knowledge

Animal behavior is the joint effect of environmental sensory impingement and the animal's emotions and knowledge of the situation. In Rolls' evolutionary theory [17], the *flexibility* of behavioral responses to reinforcing stimuli is the crucial function that explains why emotion and knowledge are so important. The essential idea is that reward and punishment specify behavioral goals, instead of particular actions, thus defining flexible routes to action. This is in contrast to stimulus-response (S-R) learning in which a particular response to a particular stimulus is learned, as described in the preceding section.

Notably, this modulation of behavior can be formulated in the same framework of reinforcement learning, and in the next two subsections we will show that pain expertise (i.e., that based on prior experiences of pain) can be a source of emotions and knowledge effectively exerting such modulation.

3.1. Emotion provides intrinsic motivation in reinforcement learning

The reinforcement learning (RL) framework is particularly appropriate for bringing learning together with what in animals one would call intrinsic motivation [2]. Despite the common perception that an RL agent's reward has to be extrinsic because the agent has a distinct input channel for reward signals, reward functions can depend on the state of the robot's internal environment, which includes remembered and learned information (see Figure 1(b)).

The dominant animal's motivation at any given time influences the arousal, strength, and direction of behavior. Similarly, each robot motivation may have a quantitative intensity and the motivation with the current highest intensity will modulate the robot reaction to stimuli. Where could robot motivation come from? Emotions are certainly motivating; for example, fear learned by stimulus-pain association provides the motivation to avoid noxious stimuli.

Rumbell et al. [19] review emotion mechanisms that have been used in artificial agents to improve action selection. Many include arousal and pleasure/pain components, to which the synthetic forces approach adds a clarity/confusion one. Emotion modulates behavior by using arousal to focus attention and associating pleasure/pain to goals, i.e., courses of action that should be pursued or avoided. Confusion tends to increase pain, and clarity tends to increase pleasure, ensuring that the agent will seek situations where its expectations are confirmed.

But is robot emotion (e.g., fear of pain) possible? Currently, there are two main programs of scientific research investigating this question: psychological behaviorism and role functionalism. Psychological behaviorism assumes that emotion behavior can be explained in a robot without reference to subjective or mental states. On the contrary, a role functionalist defines pain as "a subjective qualitative state that is reliably caused by noxious stimulation, to cause the desire to make it stop, to cause distraction regarding concurrent projects or plans and their completion, and to cause changes in preferences among alternative states of affairs" [27]. According to role-functionalism, only beings with subjective qualitative states that fulfill these causal roles can be in pain.

Following this view, Parisi and Petrosino [16] equipped robots with an emotional circuit and devised five experiments, one of them exploring robot pain. Their claim was that, for robots to "experience" pain, they need to have more than one dispositional motivation; and competition between the motivations must result in an overriding aversion, or some disposition to avoid if the result is to be painfulness. In the experiment, the emotional circuit succeeded in arousing a compelling aversion, despite concurrent hunger. Thus, in the terms of role-functionalism, Parisi and Petrosino's robots are considered to be in pain because it is possible to describe the functional role that pain states have in their behavior and to identify the specific part of the neural network causing their behavior that makes pain states possible.

3.2. Reinforcement learning to plan

To go beyond sensorimotor adaptation and emotionally-modulated reactive behavior, planning capabilities are needed and, in order to develop them, a cognitive architecture [28] should be able to progressively learn an action model from experiences and rehearse hypothetical future scenarios on that model so as to determine the best course of action. Moreover, such action model must be probabilistic to account for noise in perceptions and uncertainties in action outcomes.

Relational reinforcement learning [23] has been developed towards this aim. By enhancing RL with relational representations of states and actions (i.e., explicitly encoding relations in a symbolic data structure), the knowledge acquired by the RL agent can be generalized across states and transferred across tasks [13].

Some robot actions may be irreversible leading to unrecoverable failures (e.g., damaging some robot element, breaking an object or losing a tool). On the contrary, a planner can always backtrack from a dead-end so as to try to find an alternative sequence of actions to reach the goal. Thus, prior experiences of pain can be very helpful for safety-aware planning, i.e., to generate plans that preserve robot integrity while pursuing task accomplishment.

We have proposed a relational RL method that allows a robot to reason about dead-ends and their causes. If a plan might lead to a dead-end (e.g., one that potentially includes a painful action effect), the robot tries to find an alternative safe plan and, if not found, it asks a teacher whether the risky action should be executed. This method permits learning safe policies as well as minimizing unrecoverable errors during the learning process, and it has been validated on a robot learning to clear a table [12].

4. Conclusions and future prospects

Given the scarce bibliography dealing explicitly with robot pain, this chapter has tried to enrich its review with related research works about robot behaviors and capacities in which pain could play a role. It has been shown that all such roles —ranging from punishment to intrinsic motivation and planning knowledge— can be formulated within the unified framework of reinforcement learning.

We foresee that *self-knowledge* will be the key ingredient to significantly increase robot autonomy in the years to come and that pain could serve to acquire such knowledge. Several degrees of self-knowledge could be distinguished progressively leading to more complex functionalities. The simplest body schema consisting of a parameterized kinematic/dynamic model has already been incorporated into robots through sensorimotor learning. More elaborate self-models would include a body image precisely delimiting the robot boundaries acquired through exploratory actions, leading to the distinction between a body self and others [15], from which the notions of body ownership and agency could develop. These would ground the foundations for robots to construct a model of their own physical/cognitive abilities which, for instance, could allow them to ask for help whenever a task goes beyond their capabilities or either explore new actions [1] and try to acquire the required skills if no helper is around. Finally, some authors hypothesize that the most sophisticated internal models would lead to consciousness [8]. Whether this intriguing, largely unknown feature of the human brain could serve a functional purpose in robots constitutes an amazing challenge for future technological research, while at the same time posing decisive questions that promise to ignite an exciting moral and ethical debate.

REFERENCES

[1] Agostini A., Torras C., Wörgötter F. (2015) An Efficient and Interactive Decision-Making Framework for Robotic Applications. *Artificial Intelligence*, Special Issue 'AI and Robotics'. Available online: http://dx.doi.org/10.1016/j.artint.2015.04.004

[2] Barto A.G. (2013) Intrinsic motivation in reinforcement learning. In G. Baldassarre and M. Mirolli (eds.): *Intrinsically motivated Learning in Natural and Artificial Systems*, pp. 17-47, Springer-Verlag Heidelberg Berlin.

[3] Cardinali L., Frassinetti F., Brozzoli C., Urquizar C., Roy A.C., Farnè A. (2009) Tool-use induces morphological updating of the body schema. *Current Biology*, *19*(12), R478-R479.

[4] de Vignemont F. (2015) Pain and the spatial boundaries of the bodily self. In this volume.

[5] Fellous J.M., Arbib M.A. (Eds.) (2005) "*Who needs emotions? The brain meets the robot".* Oxford, UK: Oxford University Press.

[6] Haselager W.F.G., Broens M., Gonzalez M.E.Q. (2011) The importance of sensing one's movements in the world for the sense of personal identity. *Rivista Internazionale di Filosofia e Psicologia*, *3*(1), 1-11.

[7] Hoffmann M., Gravato H., Hernandez A., Sumioka H., Lungarella M., Pfeifer R. (2010). "Body schema in robotics: A review". *IEEE Trans. on Autonomous Mental Development*, 2(4), 304-324.

[8] Holland O., Goodman, R. (2003) Robots with Internal Models: A Route to Machine Consciousness? *Journal of Consciousness Studies*, 10(4-5): 77-109.

[9] Kober J., Bagnell D., Peters J. (2013) Reinforcement Learning in Robotics: A Survey. *International Journal of Robotics Research*, 32(11): 1238-1274.

[10] Koos S., Cully A., Mouret J.B. (2013) Fast damage recovery in robotics with the t-resilience algorithm. *International Journal of Robotics Research*, 32(14): 1700-1723.

[11] Maravita A., Iriki A. (2004) Tools for the body (schema). *Trends in Cognitive Sciences*, 8(2): 79-86.

[12] Martínez D., Alenyà G., Torras C. (2014) Finding safe policies in model-based active learning. *IROS Workshop on Machine Learning in Planning and Control of Robot Motion*, Chicago, pp. 1-6.

[13] Martínez D., Alenyà G., Torras C. (2015) Relational Reinforcement Learning with Guided Demonstrations. *Artificial Intelligence*, Special Issue 'AI and Robotics'. Available online: http://dx.doi.org/10.1016/j.artint.2015.02.006

[14] Millán J.R., Torras C. (1999) Learning sensor-based navigation. In K. Morik, M. Kaiser and V. Klingspor (eds.) *Making Robots Smarter: Combining Sensing and Action through Robot Learning*, pp. 85-108, Kluwer Academic Publishers: Boston, MA.

[15] Oudeyer P-Y., Kaplan F., Hafner V.V. (2007) Intrinsic motivation systems for autonomous mental development. *IEEE Trans. on Evolutionary Computation*, 11(2), 265-286.

[16] Parisi D., Petrosino G. (2010) Robots that have emotions. *Adaptive Behavior*, 18(6): 453-469.

[17] Rolls E.T. (2013) *Emotion and decision-making explained*. Oxford University Press.

[18] Ruiz de Angulo V., Torras C. (1997) Self-calibration of a space robot. *IEEE Transactions on Neural Networks*, 8(4): 951-963.

[19] Rumbell T., Barnden J., Denham S., Wennekers T. (2012) "Emotions in autonomous agents: comparative analysis of mechanisms and functions". *Autonomous Agents and Multi-Agent Systems*, 25(1), 1-45.

[20] Simon H.A. (1969) *The Sciences of the Artificial*. MIT Press, Cambridge, MA.

[21] Stamenov M.I. (2005) Body schema, body image, and mirror neurons. In H. De Preester and V. Knockaert (eds.) *Body Image and Body Schema: Interdisciplinary perspectives on the body*, vol. 62, John Benjamins Publishing, pp. 21-43.

[22] Sutton R.S., Barto A.G. (1998) *Reinforcement Learning: An Introduction*. MIT, Cambridge. https://www.dropbox.com/s/b3psxv2r0ccmf80/book2015oct.pdf (2nd ed.).

[23] Tadepalli P., Givan R., Driessens K. (2004) Relational reinforcement learning: An overview. *ICML Workshop on Relational Reinforcement Learning*, pp. 1-9.

[24] Torras C. (1995) Robot adaptivity. In Steels L. (ed.) *The Biology and Technology of Intelligent Autonomous Agents*. Springer Berlin Heidelberg, pp. 53-71.

[25] Torras C. (2003) Robot arm control. In Arbib M.A. (ed.) *Handbook of Brain Theory and Neural Networks* - 2nd edition, MIT Press: Cambridge, Massachusetts, pp. 979-983.

[26] Ulbrich S., Ruiz de Angulo V., Asfour T., Torras C., Dillmann R. (2009) Rapid learning of humanoid body schemas with kinematic Bézier maps. *9th IEEE-RAS International Conference on Humanoid Robots*, Paris, pp. 431-438.

[27] van Rysewyk S. (2013) Robot pain. *International Journal of Synthetic Emotions*, 4(2), 22-33.

[28] Vernon D., Metta G., Sandini G. (2007) A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation*, 11(2), 151-180.