A Flexible Component-Based Robot Control Architecture for Hormonal Modulation of Behaviour and Affect*

Luke Hickton, Matthew Lewis, Lola Cañamero

Embodied Emotion, Cognition and (Inter-) Action Lab School of Computer Science, University of Hertfordshire College Lane, Hatfield, Herts AL10 9AB, UK

Abstract

In this paper we present the foundations of an architecture that will support the wider context of our work, which is to explore the link between affect, perception and behaviour from an embodied perspective and assess their relevance to Human Robot Interaction (HRI). Our approach builds upon existing affect-based architectures by combining artificial hormones with discrete abstract components that are designed with the explicit consideration of influencing, and being receptive to, the wider affective state of the robot.

1 Introduction

The ability of embodied agents to integrate within a human-centric environment may depend upon their capacity to respond to affective and behavioural cues: attainment of their goals, and possibly even their survival, could be contingent on their capacity to interpret and convey emotion.

Affect-based systems can provide a practicable mechanism of managing internal resources and conflicting goals. They can also facilitate the expression of needs so as to elicit appropriate emotional responses and empathy from human observers, which is important in the context of Human Robot Interaction (HRI).

Research on affect has traditionally been focussed on one of two directions: works that explore emotional traits aimed at facilitating interaction with humans, versus those that focus on the adaptive mechanisms that are advantageous for survival. Some authors (e.g. [1, 2, 3]) suggest that the former approach tends to model 'shallow' or 'superficial' aspects of emotions unless combined with the latter. This view is consistent with the perspective that emotions should be grounded in the agent's architecture and internal value system, rather than

*Published in the proceedings of 18th Towards Autonomous Robotic Systems (TAROS) Conference 19th-21st July 2017, University Of Surrey, United Kingdom. The final publication is available at Springer via 'https://doi.org/10.1007/978-3-319-64107-2_36'

modelled from the point of view of an observer [1, 4]: a key tenet of the Animat approach [5].

Damasio suggests that emotion is a vestigial mechanism for maintaining homeostasis [6], hence modelling homeostatic variables and hormone production can provide a compelling approach to the grounding problem [7]. Brooks' summary of Kravitz's work on hormonal responses in lobsters [8], defines eight distinct principles in his computation model of hormones [9]. This paper focuses on the fourth of these: effects of hormones on 'sensory elements...higher processing centers, and motor or hormonal output systems'.

Brooks' implementation [9] of Kravitz's work controlled sensory elements, higher processing centers, and motor output systems via behaviours, which are in turn influenced by hormones. Subsequent works [2, 10, 11] have tended to take a similar position, focusing on gross level systems like the Action Selection Mechanism (ASM) to determine behaviour, rather than allowing it to occur as a consequence of modulating finer-grained elements of sensory, cognitive or motor function. This is analogous to mechanisms like pupil dilation in mammals, occurring as part of the 'fight or flight' stress response triggered by the sympathetic branch of the nervous system. Pupil dilation has the benefit of facilitating predator detection, which is a prerequisite of any behaviour intended to increase the likelihood of evading capture. Hence hormone induces behaviour as result of altering the sensory experience, rather than by changing the response to it.

This distinction can be illustrated in terms of behaviour-based robotics by a scenario using stress hormones to modulate response to a predator. On the one hand, an approach that modulates perception at the gross level, such as the one adopted in [12], uses hormonal modulation of exteroception (i.e. perception of external stimuli) to intensify the predator's perceived proximity, and increase the likelihood that the ASM will select a behaviour consistent with predator evasion. A similar method uses hormonal modulation of interoception (i.e. hormonal modulation of internal stimuli) to amplify the internal perception of tissue damage, causing the ASM to select a behaviour tailored to the detection of potential threats in order to prevent further injury. In both cases, the behaviour of the robot changes, whilst the behaviour of the sensor remains the same.

This contrasts with models that create direct associations between hormone levels and the functional properties of individual sensory elements, such as their range or update frequency [13]. Following this other approach, the presence of stress in the above example could increase the energy employed by the sensors for the purposes of threat detection. Whilst this could be a useful adaptive mechanism in isolation, it becomes more powerful when several sensors are combined with 'higher processing centres' [8] to form simple sensory systems, as we seek to demonstrate later in this paper. These sensory systems can provide more properties that can be modulated by affect, and provide a degree of autonomy that more closely resembles the reflexive mechanisms found in biological systems, such as saccadic responses in the mammalian vision system. Furthermore, by modelling these characteristics, and grounding their operation more firmly in the affective context of the robot, we postulate that the resulting behaviour driven by these systems will appear more natural and expressive to human observers.

Whilst this discussion has focussed predominantly on sensory systems, other functional units that could benefit from being grounded to the robot's affective state include motor output systems [8]. Kinesic properties such as posture, quality of movement and motion dynamics are all good communicators of affect in mammals [14]. Anger, for example, is usually expressed by large, jerky movements coupled with an erect stance whilst sadness is often characterised by collapsed posture coupled with small slow movements [15]. This kinesic colouring can determine how a core behaviour is interpreted by an observer: avoidance behaviour combined with 'angry' kinesics could give the appearance of a creature that is seeking to avoid confrontation, but which is ultimately prepared to fight. Conversely avoidance combined with 'sadness' might convey resignation or exhaustion.

We have described the benefits of separating the coarse-grained elements that remain the domain of the primary ASM, the 'what', from the finer-grained nuances which are critical to the 'how'. The former relates to the satisfaction of motivations via appropriate action, whilst the later pertains to regulation of internal systems to maximise their relevance to the robot's present context. In this paper we propose how sensory systems, motor output and other 'higher processing centers' [8] could be modelled via a decentralised architecture consisting of loosely coupled components that use hormones to coordinate discrete aspects of their operation. In doing so we hope to describe the benefits from both an adaptive viability and an HRI perspective.

After briefly outlining our architecture, we will consider the key properties of hormones and how they can be represented using components and sockets. An example component, the Perceptual Memory Map, will be introduced followed by an illustration of how elements of its operation can be coordinated via sockets. We will then suggest how these elements could be integrated within a wider Motivational Action Control Architecture before concluding with a summary of the key benefits of this approach and an outline of future work.

2 Architecture

This section presents a robot architecture that has been designed to model some of the key properties of hormone interaction upon the wider system. Hormones have the ability to target multiple areas of the nervous system in different ways and for different durations. In addition to the level of hormone present, salient factors include the site's sensitivity to the hormone, presence of inhibitors and duration of exposure [12, 16, 17]. We have attempted to reproduce these features by combining functional units, representing areas of the nervous system, with interfaces that enable their properties to be connected to a hormone source via a weighted link. We have labelled the former constructs 'components' and the latter 'sockets'.

2.1 Components

Our first component is the Perceptual Memory Map (PMM), which provides a residual memory of captured sensory data, coupled with a confidence estimate of how reliable that data is, as explained in detail below. It is designed to model cognitive phenomena such as attention and surprise, and models the area in the immediate surroundings of the robot by dividing it into a variable number of segments. When a new sensor value is obtained, the segment closest to its azimuth is updated, as shown by Fig. 1a below. Following movement of the agent, the recorded sensor values and their corresponding angles are mapped into two-dimensional space and transformed relative to the robot's new position, shown by Fig. 1b and Fig. 1c respectively. If two samples share the same angle, the smallest is retained. A potential fields mechanism is used to calculate a desired movement vector, which is then interpreted by the movement controller, as illustrated by Fig. 1d. The model can also output a vector indicating the most advantageous movement in terms of sensory data acquisition, which the action selection mechanism can integrate with the movement vector, depending on the holonomic constraints of the robot.

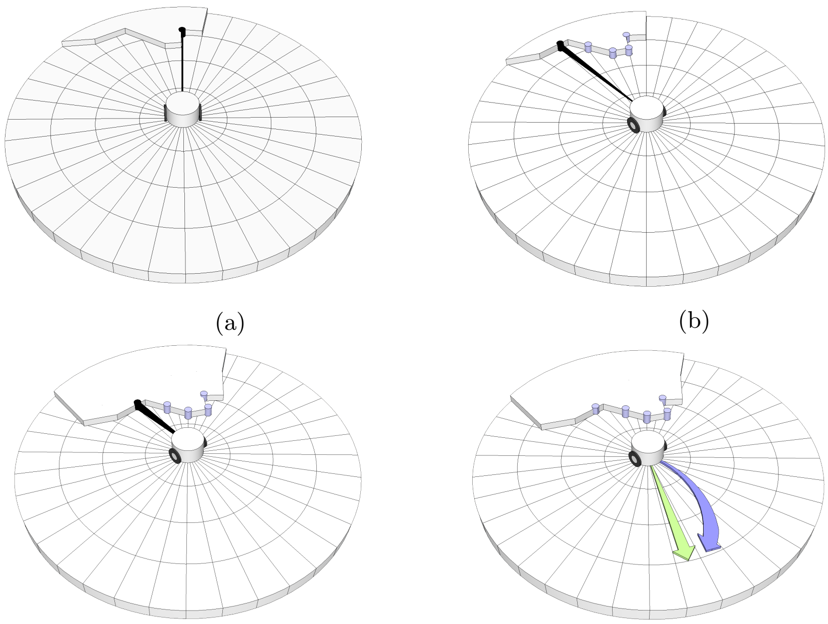

(a) (b) (c) (d)

Figure 1: Diagrams showing the state of the Perceptual Memory Map after: (a) a single proximity sensor captures distance in real-time, (b) the robot rotates on the spot by 30 degrees, (c) the robot has moved forward 10 mm and (d) the resultant movement vector is calculated.

The transformation of sensory data is contingent on the orientation and position of the robot, which we obtain via differential drive forward kinematic equations, using live motor speeds sampled at 100 ms intervals. This open-loop approach fails to take account of acceleration or deceleration occurring between the sample periods or external physical factors, such as loss of traction, but these cumulative errors are of little concern since the captured sensor values decay over time. This decay represents loss of confidence due to environmental changes that could have occurred without the agent's knowledge, which O'Regan labels insubordinateness [18], and also limits to the robot's ability to accurately track its location over extended time periods. When a value is first captured by a sensor the recorded confidence is 100%, which decreases as the sensor is moved away. If the confidence level decays to zero, the value is removed. Similarly,

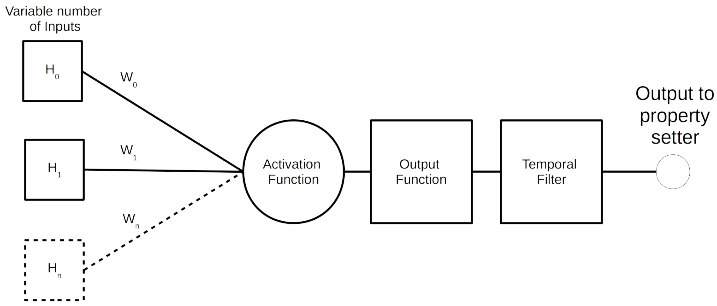

Figure 2: Diagram showing the structure of an input socket, connected to a variable number of hormone sources by weighted links.

it is also removed if the robot is more than a given distance from the mapped value. Later experiments will seek to control the decay, range and resolution of the model via hormonal control.

Our model was inspired by the findings of studies on spatial mapping in rats and path integration in cockroaches and other insects. Rats were found to have specialised cells in the parahippocampal cortex: head direction cells that fire when the animal faces a specific direction, grid cells that map distances and border cells which represent proximity to boundaries [19]. Path integration research on insects revealed how proprioceptive sensation is used to measure distance travelled and orientation relative to their environment [20, 21] which is consistent with the approach we have taken. The PMM is flexible enough to be useful in a number of contexts. We have applied it to the whole agent, but it could equally be utilised at a more localised level, for example simulating saccadic movement in an ocular sensory system.

2.2 Sockets

Sockets are software constructs that enable weighted connections to be made between hormones and the discrete properties components make available for hormones modulation. There are two types of socket: input and output sockets. Input sockets enable many hormones to modulate a single property, whereas output sockets enable a single property to affect a number of hormones. Figure 2 illustrates the structure of the input socket. In common with a Perceptron [22] the input socket incorporates an activation function and an output function. An additional 'temporal filter' inhibits the output if it has been active for too long. Unlike a Perceptron, our activation function calculates the average of the weighted inputs, rather than the sum. The inputs are constrained to values between 0 and 100, hence the output from the activation function will also be within this range. Since we are currently using a linear output function, the value passed to the component's property will be the same as the socket's activation. The temporal function will be explored in future experiments, but is currently unused. Output sockets allow components to influence hormone levels within the system via a weighted link from the property to the hormone. The inbound connection to the hormone itself is made using an input socket.

Figure 3 provides an example which illustrates how a single hormone can

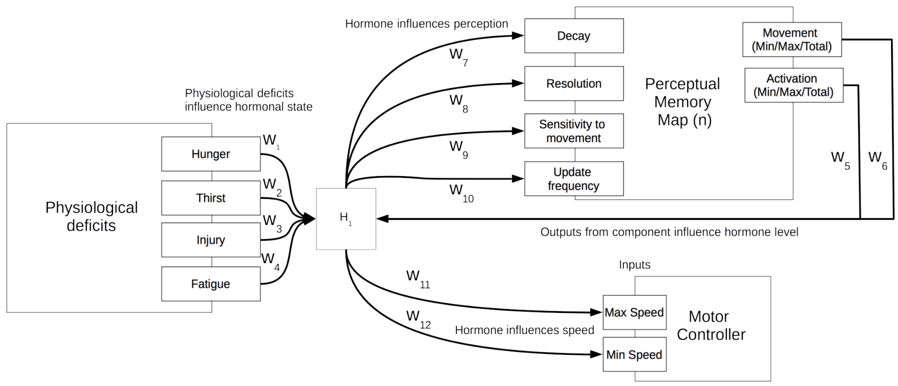

Figure 3: Diagram showing a socketed architecture in which a single hormone modulates aspects of perception and motor control.

simultaneously modulate both perception and motor function via socketed components. The two components considered are a PMM and a Motor Controller. The PMM has input sockets to set decay rate, resolution of the map, sensitivity to movement and update frequency. These are connected to weights 7, 8, 9 and 10 respectively. The Motor Controller's input sockets, connected to weights 11 and 12, enable minimum and maximum speeds to be changed. Therefore specifying positive values for weights 7 to 11 would result in greater alertness coupled with faster movement speeds, simulating some of the aspects of hormones such as Epinephrine, which trigger physiological arousal. The PMM also has an output socket that connects the presence of movement to the hormone level via weight 6.

Providing positive values for this weight would increase the levels of hormone whenever movement was detected, simulating the roles of the endocrine system and hypothalamus in mammals. This example shows the grounded physiological consequences of affect in precipitating the allocation of internal resources in response to environmental challenges, and behaviour that an external observer can interpret as an emotional response, in this case anxiety. It also shows how the architecture facilitates the coordinated change of individual subsystems, which collectively result in a behavioural response relevant to the agent's affective state.

2.3 Integrating the Perceptual Memory Map as Part of a Motivational Action Selection Control Architecture

Figure 4 illustrates how the perceptual map component has been integrated into a motivational control system, based on our group's longstanding approach [23, 12, 24]. Our architecture has been designed to enable the agent to interact in the environment shown in Fig. 5.

Besides the robot, initially a Khepera 3, environmental features include large blocks and smaller pillar-like objects, designed to require careful detection for the robot to avoid them successfully. Predator and prey are represented by red and green robots which slowly move around the arena, attracted to and repelled by the agent respectively. A blue circle of 10cm in diameter is affixed to the

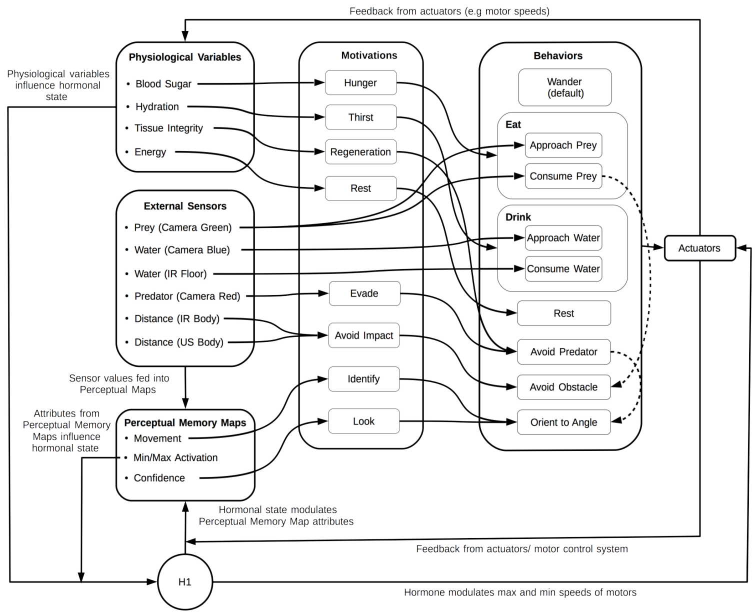

Figure 4: Diagram showing the Perceptual Memory Map integrated in an Extended Motivational Action Control System.

floor of the arena to represent a static source of hydration. Maintenance of homeostasis is dependent upon four physiological variables: blood sugar, hydration, tissue integrity and energy. Blood sugar is replenished by collision with a mobile prey robot; hydration is restored by moving over a blue circle in the arena; energy is depleted by movement and replenished by inactivity and tissue damage occurs as a result of collisions with static obstacles or mobile predators. Deficits in the homeostatic variables drive motivations to rectify the imbalance, which in turn weight the ASM towards appropriate behaviours. Behaviours will not be selected without appropriate motivations, but additional sensory stimuli may also be required. For example, water will not be consumed unless there is a motivation to drink coupled with the presence of water, which is indicated by stimulation of the floor facing IR sensor. Actions that are dependent on both motivations and sensory cues are represented by shaded boxes on the diagram.

Predators, prey and water are detected via a front-facing camera which captures levels of red, green and blue in a 60 degree field of view. Obstacles are sensed via a ring of infra-red (IR) and ultrasound emitters and objects less than 20mm away from the robot will generally cause damage unless the the prey sensor is highly stimulated. In this case, damage is suppressed for forward-facing sensors and blood sugar is replenished instead. Water is consumed via stimulation of a floor facing IR sensor. The proximity sensors are also associated with the PMM, which are able to pass vectors and other contextual information into the ASM. Output properties of the PMM include movement detection, activation and confidence. Movement will trigger a motivation to identify the source with the main camera, whilst low confidence in the validity of its internal representation will trigger a motivation to look around. The total activation

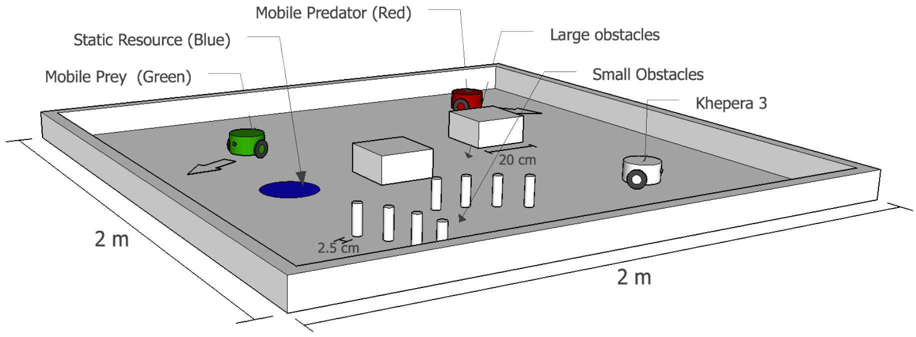

Figure 5: Diagram showing the environment the agent was designed to occupy.

of the PMM is used to influence hormone levels within the system, in this case simulating fear of being trapped, and hormone level is also increased by physiological deficits.

The input sockets are not shown on the diagram for the sake of clarity, since they have been discussed in previous sections, but elevated levels of hormone increase motor speed, heighten sensitivity to movement and reduce overall confidence, creating behaviour that might equate to anxiety. Note that the implementation details of the ASM are not considered to be important, since the architecture is compatible with most action selection approaches, including winner-takes-all and voting-based methods. However, our implementation uses a winner-takes-all approach, and the inbound connections to each behaviour are assumed to be directly proportional to the probability of it being selected as the dominant activity.

Coupling the PMM tightly with the hormone facilitates the positioning of sensors with the aim of maximising their acuity and relevance to the agent's affective context, allowing it to actively drive behaviour that meets the current requirements of the sensory system. The ASM helps coordinate this behaviour with competing demands from other sensory systems. We believe this mechanism is scalable, and can support many other components that will collectively produce more behavioural properties that human observers will be able to identify with those of other creatures.

3 Discussion

This paper has described an architecture intended to facilitate subtle forms of expression and adaptive behaviour. It has also suggested how the properties of sensory perception, higher reasoning and motor control can be represented by loosely coupled components coordinated via hormones. We believe this granular approach provides the following benefits: a grounded link between affect and behaviour, an approach that can be implemented using a variety of paradigms, and an architecture that can be applied to different morphologies of robot with only minor adjustment.

In terms of grounding, we have endeavoured to ensure our model of affect directly influences aspects of perception, cognitive function and motor activity. This granular, component-driven focus reflects the findings of neuroscience research suggesting that different sensory aspects, such as colour and edge detection, are spatially separated in the brain, even within the same sensory systems [25]. Regarding flexibility of implementation, we represent components via programatic constructs but other paradigms, such as Artificial Neural Networks, could have been adopted instead [26, 27]. A benefit of our design is that different approaches can also be used interchangeably within the same architecture, which enables broad components to be used as temporary proxies until they can be decomposed into smaller units of functionality. Finally, we contend that our architecture promotes flexibility, since it facilitates encapsulation of elements that are specific to the morphology of the platform, such as sensor input and motor output. If, for example, the architecture is migrated to a platform that has additional motor capabilities, these additional properties can easily be associated with the affect mechanism and coordinated via this secondary control system without making any fundamental changes to the behaviours or ASM.

We believe this model could help simple situated robots to respond to the affective and behavioural cues of humans and other actors and facilitate communication of their needs in a way that mirrors the familiar forms of non-verbal communication which are typical of many mammals. However, we are still in the preliminary stages of empirical testing and are currently carrying out experiments to assess the adaptive value of the PMM in different environments and with different sensory configurations.

4 Acknowledgements

Luke Hickton is supported by a PhD studentship of the University of Hertfordshire.

References

[1] Cañamero, L.: Building emotional artifacts in social worlds: Challenges and Perspectives. Proceedings of 2001 AAAI Fall Symposium in Emotional and Intelligent II: The Tangled Knot of Social Cognition pp. 22–30 (2001)

[2] Cañamero, L.: Emotion understanding from the perspective of autonomous robots research. Neural Networks 18(4), 445–455 (2005)

[3] Herrera Perez, C., Sanchez Escribano, G., Sanz, R.: The morphofunctional approach to emotion modelling in robotics. Adaptive Behavior 20(5), 388–404 (2012)

[4] Pfeifer, R.: Building "Fungus Eaters": Design Principles of Autonomous Agents. In Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior SAB96 (From Animals to Animats) pp. 3– 12 (1996)

[5] Wilson, S.: The Animat Path to AI. In: From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour. vol. 1, pp. 15–21. Cambridge, MA: MIT Press (1991)

[6] Damasio, A.R.: The Feeling of What Happens: Body and Emotion in the Making of Consciousness (1999)

[7] Harnad, S.: The Symbol Grounding Problem. Physica D 42, 335–346 (1990)

[8] Kravitz, E.A.: Hormonal control of behavior: Amines as gain-setting elements that bias behavioral output in lobsters. Integrative and Comparative Biology 30(3), 595–608 (1990)

[9] Brooks, R.A.: Integrated systems based on behaviors. ACM SIGART Bulletin 2(4), 46–50 (1991)

[10] Krichmar, J.L.: A neurorobotic platform to test the influence of neuromodulatory signaling on anxious and curious behavior. Frontiers in Neurorobotics 7(FEB), 1–17 (2013)

[11] Lowe, R., Kiryazov, K.: Utilizing Emotions in Autonomous Robots: An Enactive Approach, pp. 76–98. Springer International Publishing, Cham (2014)

[12] Cañamero, L., Avila-García, O.: A bottom-up investigation of emotional modulation in competitive scenarios. In: Paiva, Ana C. R.and Prada, Rui and Picard, R.W. (ed.) Affective Computing and Intelligent Interaction: Second International Conference, ACII 2007 Lisbon, Portugal, September 12-14, 2007 Proceedings. pp. 398–409. Springer Berlin Heidelberg, Berlin (2007)

[13] Lones, J., Lewis, M., Cañamero, L.: From Sensorimotor Experiences to Cognitive Development: Investigating the Influence of Experiential Diversity on the Development of an Epigenetic Robot. Frontiers in Robotics and AI 3(August) (2016)

[14] Lhommet, M., Marsella, S.: Expressing Emotion Through Posture and Gesture. The Oxford Handbook of affective computing (Gratiolet 1865), 273–285 (2015)

[15] Beck, A., Canamero, L., Bard, K.A.: Towards an Affect Space for robots to display emotional body language. 19th IEEE International Symposium on Robot and Human Interactive Communication Principe pp. 464–469 (2010)

[16] Gadanho, S.C., Hallam, J.: Robot Learning Driven by Emotions. Adaptive Behavior 9(1), 42–64 (2001)

[17] Levitan, I.B., Kaczmaret, L.K.: The Neuron: Cell and Molecular Biology. Oxford University Press, 3rd editio edn. (2002)

[18] O'Regan, J.K., Noe, A.: What it is like to see: A sensorimotor theory of perceptual experience. Synthese 129(1), 79–103 (2001)

[19] Derdikman, D., Moser, E.I.: A Manifold of Spatial Maps in the Brain. Space, Time and Number in the Brain 14(12), 41–57 (2011)

[20] Durier, V., Rivault, C.: Path integration in cockroach larvae, Blattella germanica (L.)(insect: Dictyoptera): Direction and distance estimation. Learning & Behavior 27(1), 108–118 (1999)

[21] Dyer, F.C., Dickinson, J.A.: Sun-Compass Learning in Insects: Representation in a Simple Mind. Current Directions in Psychological Science 5(3), 67–72 (1996)

[22] Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. Psychological review 65(6), 386–408 (1958)

[23] Cañamero, L.: Modelling motivation and emotions as a basis for intelligent behavior. In: first intl. conf. on autonomous agents (AGENTS'97). pp. 148–155. ACM Press, New York (1997)

[24] Lewis, M., Canamero, L.: Hedonic quality or reward? A study of basic pleasure in homeostasis and decision making of a motivated autonomous robot. Adaptive Behavior 24(5), 267–291 (2016)

[25] Fellous, J.: The Neuromodulatory Basis of Emotion. The Neuroscientist 5(3), 283–294 (1999)

[26] French, R.L.B., Cañamero, L.: Introducing Neuromodulation to a Braitenberg Vehicle. Proceedings - IEEE International Conference on Robotics and Automation 2005(April), 4188–4193 (2005)

[27] Husbands, P.: Evolving robot behaviours with diffusing gas networks, pp. 71–86. Springer Berlin Heidelberg, Berlin, Heidelberg (1998)