Robotics and Autonomous Systems 68 (2015) 12–20

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

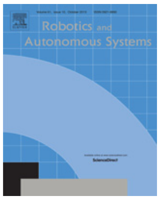

Intrinsically motivated learning systems based on biologically-inspired novelty detection

Y. Gatsoulis [a],*, T.M. McGinnity [b],[1]

[a] School of Computing, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, UK
[b] Intelligent Systems Research Centre, University of Ulster, Northland Road, Derry BT48 7JL, UK

HIGHLIGHTS

- We proposed novelty detection as intrinsic motivation for autonomous learning.
- Novelty detection model is bio-inspired by the behavioural phenomenon of habituation.
- We demonstrated the effectiveness of the model in a real-world experiment.
- We showed how habituation is used as an intrinsic motivation for robot learning.

ARTICLE INFO

Article history:
Received 5 August 2014
Received in revised form
31 January 2015
Accepted 16 February 2015
Available online 21 February 2015

Keywords:
Intrinsic motivations
Novelty detection
Habituation
Unsupervised learning
Bio-inspired robots

ABSTRACT

Intrinsic motivations play an important role in human learning, particularly in the early stages of childhood development, and ideas from this research field have influenced robotic learning and adaptability. In this paper we investigate one specific type of intrinsic motivation, that of novelty detection and we discuss the reasons that make it a powerful facility for continuous learning. We formulate and present one original type of biologically inspired novelty detection architecture and implement it on a robotic system engaged in a perceptual classification task. The results of real-world robot experiments we conducted show how this original architecture conforms to behavioural observations and demonstrate its effectiveness in terms of focusing the system's attention in areas that are potential for effective learning.

1. Introduction

There is a long history of research in motivational and drive theories for humans [1–4]. Of particular interest is the concept of intrinsic motivation, which if transferred appropriately to the domain of robotics can lead to the research and development of effective unsupervised continuous learning machines [4]. Intrinsic motivations are the drive to engage in an activity because it is "inherently enjoyable" without the need for an explicit reward, in contrast to extrinsic motivation which initiates engagement in an activity because of some specific reward outcome [3,5].

It is argued that intrinsically motivated behaviour is essential for an organism to gain competences and that there is ample evidence that the opportunity to explore a novel environment can itself act as a reward, as argued by White in [1]. His hypothesis is that not only is exploration incited by novelty, but also that manipulation, or just activity itself can be rewarding; and in terms of function, intrinsic motivations are drives that guide learning of skills and knowledge. Another proposal is that of cognitive dissonance [6], in which motivation is a process directed to minimise dissonance between the internal model of the environment and perceptions. The concept of optimal incongruity is suggested in [7] which argues that "interesting" stimuli are those that differ from "standard" stimuli, i.e. novel or unexpected stimuli.

Within neuroscience, one of the most articulated and empirically well-supported theories [8–11] states that intrinsic motivations are related to the function of phasic dopamine (DA) neuron activity caused by the activation of the superior colliculus located in the midbrain. This theory claims that DA represents a learning signal that is generated by the colliculus when the organism perceives the sudden occurrence of an unexpected sensory event (e.g. a sudden light going on) and that the response of DA neurons to the sensory event ceases after the sensory event becomes predictable, e.g., by learning the action that caused it. This is particularly true

---

* Corresponding author. Tel.: +44 113 3436589; fax: +44 113 3435468.
E-mail addresses: y.gatsoulis@leeds.ac.uk (Y. Gatsoulis), tm.mcginnity@ulster.ac.uk (T.M. McGinnity).
[1] Tel.: +44 28 71675616; fax: +44 28 71675570.

http://dx.doi.org/10.1016/j.robot.2015.02.006

in the case of visual stimuli, which is the experimental scenario in this paper. The key point related to this paper is that a lack of predictability in sensory events is a form of novelty. A second important theory on intrinsic motivations within the neuro-scientific literature has been proposed in [12,13]. This theory is focused on the hippocampus and surrounding brain areas: these have been shown to respond to novel stimuli, novel associations between familiar stimuli, or familiar stimuli in novel contexts. The activation of these brain areas causes an increase in DA neuron activity which is proposed to drive the memorisation of the novel stimuli within the same areas. As learning progresses, the response of these areas progressively attenuates. This theory is particularly important for this paper as it relates to mechanisms that might represent a possible biological equivalence with the computational novelty detection mechanisms presented herein.

While these theories have some differences, they, as well as those previously mentioned [6,1,7] all regard novelty as a form of intrinsic motivation.

In order for novelty to guide exploratory behaviour, novelty must be clearly detected first. Although there is no strict definition of novelty detection, it is widely regarded [14–16] as the process of identifying stimuli that are 'different from anything known before; new, interesting and often seeming slightly strange' [17]. A formal description of the problem of novelty detection is as follows. An agent is trained to do pattern matching on a set of examples $X = x_1, x_2, \dots, x_n$ forming an initial knowledge K. At time t an observation o is considered novel if it differs significantly from what is already known, i.e. from K, using a novelty detection filter N to identify the level of novelty and the particular parts that are novel. The observation o is used as a new training example x_k, with which the agent learns and hence expands its knowledge K.

Based on these descriptions, we approach novelty detection in robotics from the perspective that an artificial agent pays less attention to perceptions that are similar to those seen during training, but is able to highlight anything different. In this sense, novelty detection can be seen as a form of selective learning: a system has some a priori knowledge, and the novelty filter aims to highlight anything that differs from what is already known. As such, novelty detection is the ability to identify perceptions that have never or infrequently been experienced before. This constitutes an important component for the effective and long-term operation of intelligent robot systems allowing computationally-efficient, unsupervised and incremental exploration and learning of new skills and environments [18,19]. However, novelty is not the same as attention. Observations in biological organisms behaviour and studies of the superior colliculus [20] and other areas show how attention is a function of several factors including, but not exclusively, novelty.

We have designed, developed and tested a biologically inspired novelty detection architecture in experiments where a physical robot cumulatively and without supervision learns the visual representation of real world objects. The results of these experiments have shown that this original architecture conforms to behavioural observations successfully and have demonstrated its effectiveness in terms of focusing the system's attention in areas that are potential for high learning. Specifically in a visual inspection and learning task we validate the system against two hypotheses that comply with observations in biological organisms behaviour. Namely these two hypotheses are:

- The robot will always pick up and train on the most novel object that lies on a table, without any human in the loop (other than changing the objects on the table). Known objects that have not been attended for some time will recover their novelty value over time, and hence become potential candidates for the attention of the system in the future.
- The robot will lose interest quicker on previously learnt objects. This results in the system spending its training time effectively, i.e. on objects in which there is a potential for greater learning.

The biologically-inspired novelty detection architecture and the discussion of the experimental results are presented in detail in the rest of the paper. First, the concept of habituation as novelty detection is described in Section 2. Related work is presented in Section 3. Section 4 explains the novelty detection model developed and the experimental scenario. Section 5 presents results and discussion, based on practical experiments. Section 6 concludes the paper.

2. Habituation as novelty detection

One mechanism by which behaviour can be directed towards investigating novel stimuli is *habituation* which is found in biological organisms [21–23]. Habituation has a long research history [24,25], and has recently regained significant attention both from a biological and psychological perspective [26,27], as well as in machine learning [28].

Habituation is a type of non-associative learning exploited to describe the behavioural phenomenon of decreased responsiveness of a cognitive organism to a recently and frequently presented stimulus. In other words habituation is a measure of novelty of a perceived stimulus depending on the familiarity of the agent with this stimulus over a duration of time; precisely, an agent has a low level of habituation over a novel a stimulus. In the rest of this paper the terms habituation and novelty detection are used to refer to this idea.

A landmark paper [24] defined habituation and described a number of its characteristics. The following are of particular interest to robotics:

- "Given that a particular stimulus elicits a response, repeated applications of the stimulus result in decreased response (habituation). The decrease is usually a negative exponential function of the number of stimulus presentations". Apart from this being the classic definition of habituation, it also suggests the computational form of habituation.
- "If the stimulus is withheld, the response tends to recover over time". The recovery of the habituation level is an important characteristic as it allows revisiting previously learnt stimuli, regardless of the level of how well these are known, which might lead to the discovery of new information about them that was not initially observable.
- "The weaker the stimulus, the more rapid and/or more pronounced is habituation. Strong stimuli may yield no significant habituation". This is another important characteristic of habituation as, if we assume for the context of robotics that the strength of the stimulus is determined by how well it is known, this means that well known stimuli will habituate faster than unknown ones. The benefit is that the learning system will spend more time with stimuli that provide the potential for greater learning, hence increasing its competency and resources.

One strength of the habituation approach in novelty detection is that it can be used to guide exploratory behaviour in the absence of a specific goal. This might lead to learning behaviours that are not immediately useful but could be in the right circumstances in the future. For example, infants may get engaged in novel activities that do not necessarily "make sense" to an adult (e.g. putting an object from the floor to their mouth); however, they choose to engage and disengage in these activities through some kind of intrinsic motivation, which in some cases is habituated novelty detection [29,30,21]. The result of these "meaningless" activities in their developmental process can lead to the acquisition of new