

Opinion piece



Cite this article: Gordon G. 2019 Social behaviour as an emergent property of embodied curiosity: a robotics perspective. *Phil. Trans. R. Soc. B* **374**: 20180029. <http://dx.doi.org/10.1098/rstb.2018.0029>

Accepted: 14 October 2018

One contribution of 17 to a theme issue ‘From social brains to social robots: applying neurocognitive insights to human–robot interaction’.

Subject Areas:

behaviour, cognition, neuroscience

Keywords:

artificial curiosity, robots, computational models, intrinsic motivation

Author for correspondence:

Goren Gordon
e-mail: goren@gorengordon.com

Social behaviour as an emergent property of embodied curiosity: a robotics perspective

Goren Gordon

Curiosity Lab, Department of Industrial Engineering, Tel-Aviv University, Tel-Aviv, Israel

GG, 0000-0002-8351-7034

Social interaction is an extremely complex yet vital component in daily life. We present a bottom-up approach for the emergence of social behaviours from the interaction of the curiosity drive, i.e. the intrinsic motivation to learn as much as possible, and the embedding environment of an agent. Implementing artificial curiosity algorithms in robots that explore human-like environments results in the emergence of a hierarchical structure of learning and behaviour. This structure resembles the sequential emergence of behavioural patterns in human babies, culminating in social behaviours, such as face detection, tracking and attention-grabbing facial expressions. These results suggest that an embodied curiosity drive may be the progenitor of many social behaviours if satiated by a social environment.

This article is part of the theme issue ‘From social brains to social robots: applying neurocognitive insights to human–robot interaction’.

1. Introduction

Social interaction is the basis of human society. It governs relationships, our place within a community and supplies information about what happens in an environment. Human children present social behaviours very early in life, as they start to perform gaze following, joint attention and produce facial expressions that are meant to provoke caregiver reactions.

Many research fields have studied different forms of social behaviours, their emergence through life and their role in normal and abnormal development. Several synergistic research tools have been used for this purpose including the following: pure observation of the behaviours of human babies at different ages has supplied numerous insights into the sequence of emerging behaviours and deviations from them [1]; methodological construction and evaluations of mental models and their development have been developed over the years, which has also supplied profound understanding of development and interactions [2]; research into the underlying mechanisms behind these behaviours [3]; recent advances in brain imaging and measurements have supplied an unprecedented view into the workings of the brain during various activities [4].

A recent field that has entered the foray approaches the subject from a unique perspective. The field of developmental robotics combines the insights from developmental psychology, neuroscience and computer science and implements them in robots, with the goal of creating robots that mimic infant development [5,6]. A sub-field of developmental robotics is called artificial curiosity, wherein a computational model of the curiosity drive has been developed [7–10]. Artificially curious robots attempt to maximize learning of themselves and their environment by selecting actions that give them the largest amount of information. Thus, they are ‘driven’ by learning, i.e. they have intrinsic motivation to learn as much as possible. These curious robots usually employ some form of reinforcement learning, wherein their (intrinsic) rewards are proportional to their learning progress [8].

In this contribution, a bottom-up constructivist approach is explored by considering a curious embodied agent, whose sole goal is to learn as much as it can about

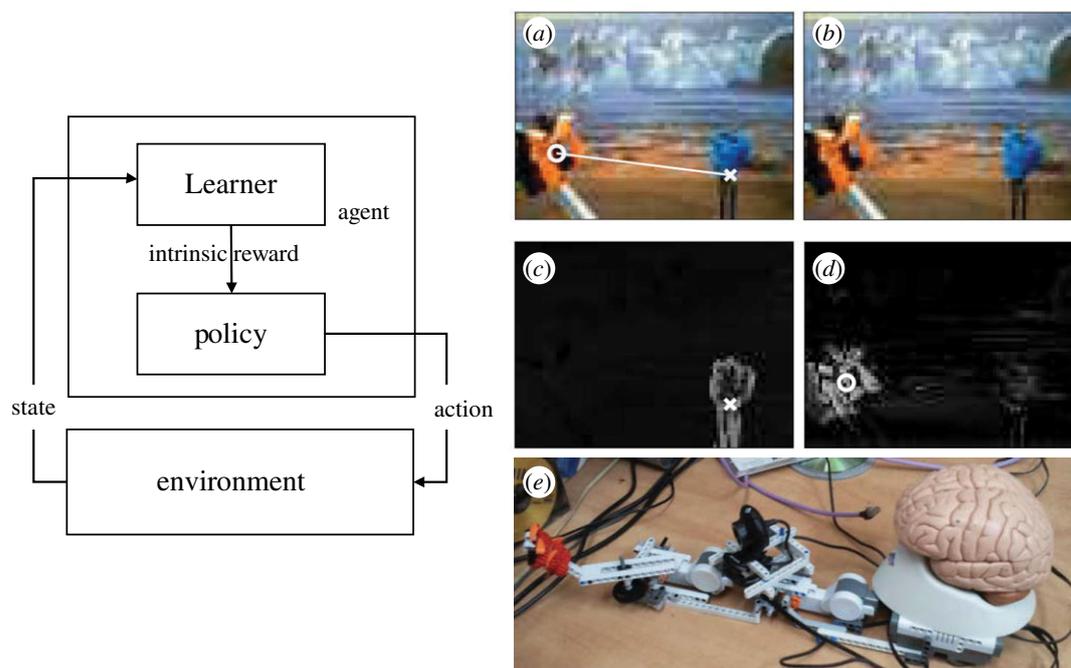


Figure 1. Left: Curiosity Loop. The agent has a Learner, which learns state–action correlations. The intrinsic rewards are proportional to the Learner’s learning progress. The Policy learns state-to-action mapping that maximizes future accumulated intrinsic rewards. Right: An embodied curious robot in non-social environments [12]. (a,b) A sequence of input images from the robot cameras. (c) Emergent motion detection Learner. (d) Emergent self-recognition Learner. (e) Robot.

itself and its environment [10–13]. We claim that there is a direct developmental path from the curiosity-based learning of pure low-level sensorimotor correlations to the emergence of social behaviours, contingent on the embedding social environment. In other words, a purely curious robot, i.e. an embodied agent implementing only an artificial curiosity algorithm, whose environment includes social agents, will inevitably learn and employ social behaviours.

Preliminary evidence is presented that demonstrates this path, based on the integration of prior studies with curious agents. The same curiosity algorithm results are shown in different developmental paths, based on the embodiment and environment as follows: in a non-social environment, the curious robot develops self-recognition and exhibits reaching movements towards moving objects [10–12], whereas a social environment results in the emergence of face detection and tracking [13], as well as generation of attention-grabbing facial expressions [14].

2. Embodied curiosity

The term ‘embodied curiosity’ refers to the application of curiosity-based computational models within an embodied, physical agent. It hints to the fact that there is a very tight interaction between the computational model and the body, i.e. sensors, motors and their physical arrangement, as well as the embedding interaction. One interaction cannot be analysed without the other, since emerging behaviours are dependent on both. In this section, these two concepts are briefly described and reviewed.

(a) Artificial curiosity

Artificial curiosity refers to the computational models wherein the goal of the agent is learning, such that it selects its action to maximize learning. The models are inspired by cognitive psychology, wherein the curiosity drive is studied and is loosely

defined as ‘an intrinsic drive to learn as much as possible’, and neuroscience, wherein recent advances have shed new light onto the mechanisms of decision-making, action selection and rewards [15,16].

On the other hand, artificial intelligence and machine learning, inspired by the biological reward system in the brain, have generated complex and detailed algorithms for action selection, one of which is reinforcement learning (RL). In this scheme, an agent’s goal is to maximize its future accumulated rewards, which are supplied by the environment. The agent learns a policy, i.e. a mapping between the state and its actions, that achieves this goal. In other words, the agent behaves in the environment and explores different policies while obtaining rewards, thus learning which policies generate the highest future accumulated rewards. Many RL algorithms have been developed; the most recent revolve around deep learning, wherein deep neural networks are used to learn the mapping between extremely complex states, e.g. images and actions [17].

Artificial curiosity is a specific extension of RL, wherein the rewards are supplied internally, hence the connection to ‘intrinsic motivation’. The goal of a curious agent is still to maximize its future accumulated rewards, but in this constellation, the rewards are internally derived from a separate learning process. A succinct representation of this is the Curiosity Loop, which is composed of a Learner and an RL module ([18,19], figure 1, left). The Learner attempts to learn some form of sensorimotor correlation, e.g. internal models, which are influenced by the actions selected by the agent. The intrinsic rewards are thus proportional to the learning progress of the Learner, i.e. the more the Learner learns to approximate the appropriate correlations, the higher the rewards. The RL component receives these intrinsic rewards and changes the policy, based on the RL algorithm. The agent ‘learns (RL) how to behave (policy) in order to learn (Learner)’, hence the Curiosity Loop. The end result of the Curiosity Loop is the convergent policy, which is the emergent behaviour from the curiosity drive.

(b) Embodiment

Embodiment refers to the actual physical body an agent has and the influence it has on other 'cognitive' processes. In the artificial curiosity context, the embodiment refers to the sensors, motors and their arrangement within the physical agent. Thus, a mobile robot with a camera will learn different sensorimotor correlations from a robotic arm with force-sensors, resulting in totally different curiosity-driven emergent behaviours.

For example, in [20], the authors used the curiosity algorithm with a simulation of the entire sound-generation track of the human body. This coupling resulted in the emergent sequence of sound generation, similar to that found in human babies. A similar algorithm was used in [21] in an elongated simple robotic platform and resulted in an emergent proximal-to-distal sequence of behaviours. Employing the artificial curiosity algorithm on a robotic finger with force sensors has resulted in the emergence of different tapping behaviours, which optimally discerned the touched surface texture [22]. Using the Curiosity Loop, Gordon *et al.* [14,19] showed that simulated curious rodents learned how to explore their physical environments on several spatial and temporal scales, resulting in exploration motor primitives that resemble their biological counterparts.

The aforementioned examples show the importance of the embodiment and the surrounding environment for the emergence of curiosity-driven behaviours. They emphasize that a single computational model, namely, artificial curiosity, can describe the emergence of a plethora of behaviours, given the specifics of the agent. We next explore in detail a developmental path that starts from pure low-level sensorimotor correlations and ends with social behaviours, depending on its environment.

3. Non-social environments

What can a curious robot learn if it is embedded in a non-social environment, i.e. it is alone with no other agents around? The simple answer is 'sensorimotor correlations', yet a more detailed account reveals much more complex aspects of this simple inquiry. Below, we follow the developmental path of two examples, namely, a humanoid robot with a sophisticated cognitive architecture [11] and a simple robot that starts from scratch and hierarchically learns new correlations, resulting in new emergent behaviours [10,12] (figure 1).

(a) Internal models and change detection

First-order sensorimotor correlations are located in the first and lowest level of the hierarchy. These are known as 'internal models' and are the correlations between sensors and their related motors, i.e. how motor movements affect the sensory input [23]. In these low-level correlations, the agent learns, for example, that being in a certain arm angle and directing the motor to move with a specific force results in a new arm angle.

However, a multi-modal account of an embodied agent results in non-trivial correlations, namely, change-detection. For example, the correlation between a camera input and a motor that moves the camera results in learning 'motion detection', based on the statistical observation that in a non-social environment when the camera moves, most pixels change, while when it does not move, most pixels do not change.

While internal models have been addressed in the machine learning community for several decades [24], in state-of-the-art robotics, such as in [11], control of the motor units is extremely

precise, not requiring *learning* of internal models. Instead, based on the robot geometry and calibrations, one can pre-program the motor forward and inverse models. Moreover, visual change detection is also pre-programmed in a rather straightforward computation of consecutive images' subtraction.

However, in an attempt to establish a true *tabula rasa*, Gordon [12] has introduced an exhaustive search of all possible sensorimotor correlations of a simple robot (figure 1e). This was done by constructing a neural network for each of the possible 2-to-1 sensorimotor correlations. In order to avoid exponential growth in the number of neural networks, and inspired by the nervous system, a pruning algorithm was used, which removed neurons that did not contribute to the learning process. Effectively, the pruning algorithm resulted in the automatic removal of networks that were trained with *uncorrelated data*, e.g. movements of one motor's correlation to the other motor's angles, since none of their neurons could contribute to the unlearnable training set. Hence, at the end of the learning process, only networks that successfully learned the 2-to-1 correlation between different channels of information from the robot, survived. The surviving networks included all internal models, i.e. forward, inverse and postdiction. Furthermore, the robot's learning has produced a neural network that learned to compute visual change detection, i.e. it received two consecutive images as input and its output was their absolute difference (figure 1c). In contrast to Saegusa *et al.* [11], this computation was *learned*, not pre-programmed.

Moreover, employing the Curiosity Loop on these low-level Learners resulted in the emergence of specific policies [10]. These policies received as intrinsic reward the neural networks' prediction error. Each learned low-level internal model has resulted in a different optimal policy that maximized its learning, e.g. moving the camera followed by not-moving it was the emergent optimal exploration behaviour for motion detection, since it resulted in alternating change and constancy in the visual field. This emergent behaviour is reminiscent of saccades of eye movements.

(b) Self-recognition

Once the agent learns to detect motion, it has another channel of information, i.e. it can now learn the second-order correlations. The agent can now learn the correlation between its visual input, its own-generated motion and motion-detection, representing the statistical observation that 'if I move my arm, and something moves in my visual field, that something is me' [11,12]. In other words, the robot autonomously learns to visually recognize itself.

In [11], a real humanoid robot was used for the experiment of actively learning self-recognition. A pre-programmed motion detection was used in combination with a visual clustering algorithm. Thus, visual blobs that moved concurrently to the arm's motion were learned to be categorized as body-parts, constituting the second-order sensorimotor correlations.

In [12], the learned lower-level internal models were used as new channels of information for another exhaustive search of correlations, using new neural networks with a pruning algorithm. This repetition of 'sprouting' of new networks followed by 'pruning' of uncorrelated channels was inspired by *exuberance* in the brain [25], in which new synapses emerge, later to be pruned by disuse.

In the second level of the hierarchy, the surviving neural networks learned the correlation between the motor

command to the arm, motion in the visual field and the actual visual input. This resulted in a neural network that received an image as an input and outputed a binary output of one for the arm and zero otherwise (figure 1d).

Moreover, the Curiosity Loop for the visual self-recognition Learner [10] converged to a unique behavioural pattern, encoded by the RL part of the Curiosity Loop. The emergent behaviour was one wherein the robot followed its own moving hand. This behaviour ensured a steady stream of informative data from the camera, as well as the arm's motor commands. This behaviour is reminiscent of infants' developmental stage of discovering their own body and revelling, i.e. getting high reward, in their own control of it.

(c) Reaching

Our curious non-social robot has discovered its own body and can use this information to learn to map the visual field (visual self-recognition) to its proprioception field (arm angles). These are the third-order sensorimotor correlations.

In [11], using active exploration of the humanoid robot's arm unit, the robot learned the arm–vision coordinate transformation. By exploring the arm's movements and position, the robot learned the proprioception correlation of its own arm and the visual coordination, as learned by the lower-level correlation of self-recognition. This arm–vision coordinate transformation enables the robot to *reach* for a visual position, i.e. move the arm so that it co-occupies the target in the visual field.

In [12], the second-order learned correlations were used as a new channel of information for the third level of the hierarchy, resulting again in an exhaustive search followed by pruning. In this level, the surviving neural networks learned to compute arm–vision coordinate transformation, by correlating the arm's proprioception coordinates with the now visually recognized arm within the visual field.

Obtaining these multi-modal mappings enabled the robot to perform reaching for a moving object. This was done by detecting the moving object by the self-learned motion detection, mapping the visual coordinates of the object to the proprioceptive field of the robot and then using its internal models to move its arm towards the moving object.

(d) Summary

To conclude this developmental path, two examples of approaches to developmental learning of non-social behaviour were presented. The complex cognitive architecture enables impressive developmental milestones in a state-of-the-art humanoid robot. By contrast, the simple non-social embodiment of the curiosity algorithm resulted in complex mapping as well as the emergence of infant-like sequences of exploration behaviours, using a bottom-up approach and an artificial 'exuberance' algorithm, which introduces new neural networks and then prunes those that cannot learn.

There are further steps along the non-social pathway, such as object detection and manipulation [26,27]. We next consider the same paradigm, but within a social environment context.

4. Social environments

What can a robot learn when it is embedded within a social environment? Moreover, what can a curious robot do to extract more information out of this social environment?

First, let us consider 'social environments'. Here, a non-verbal developmental path is followed, i.e. the agents are restricted to the visual modality. In this context, social environments consist of other agents engaged in social activities, which can be directed either at the agent or at other agents. Purely social environments are the focus, namely, people interacting and conversing in social places, e.g. apartments, restaurants, etc.

While there have been tremendous advancements in machine learning in this domain, e.g. detection of facial expression [28], social action detection [29], all have relied on supervised learning and labelled examples. By contrast, the bottom-up approach considered in this contribution presents us with computational restrictions, i.e. the agent must not be pre-programmed with any type of pattern recognition system, such as face detection, and does not have access to labelled data. Hence, the curious agent starts as a *tabula rasa*, with no prior information but the Curiosity Loop. Such a robot can only learn from informational sources in the environment (Learner) and can only try to influence the social agents (Policy) in order to supply it with more information.

(a) Faces as sources of information

Since we are interested in infant-like behaviour, it is important to consider the 'natural' visual environment of infants. It has been recently shown that the visual environment of infants is dominated by faces and then moves to hands, especially with objects [30,31]. This may hint to the special place faces hold in the social environment.

In order to explore this notion, Barkan & Gordon [13] used a sitcom (The Big Bang Theory) as a database of visual social environments, approximating a social visual scene. To overcome challenges of a moving camera, only shots in which the camera was still were taken into consideration. Within these shots, and based on the previous results, wherein an agent can learn to detect motion within a scene, changes were computed by subtracting consecutive frames. Using known face-detector algorithms [32], the hypothesis that faces dominate the information channel within a social scene was verified. In other words, visual areas that include faces incur significantly greater amounts of motion compared with the rest of a visual scene (figure 2a).

(b) Face detection and tracking

If faces are the sources of social information, then curious agents should direct their attention towards faces [33]. Barkan & Gordon [13] introduced a novel Deep Curiosity Loop architecture, wherein the Learner attempted to learn the forward model of the visual scene, i.e. predict the next frame based on the current frame and the 'localized' actions; and the RL component was composed of concatenation of nine layers of convolutional neural networks, also known as a deep Q-network (DQN) [34], which has been shown to have similar traits to the visual cortex [35]. This network received an image as input and produced a predicted value map as its output. More importantly, the DQN received the Learner's prediction error as intrinsic reward. This architecture, combined with the realization that faces incur more motion than non-faces, resulted in intrinsic reward for faces' movement, since they were the unpredictable changes. In turn, the DQN learned to associate faces with high value and to direct the 'virtual gaze' (or actions) towards faces (figure 2b).

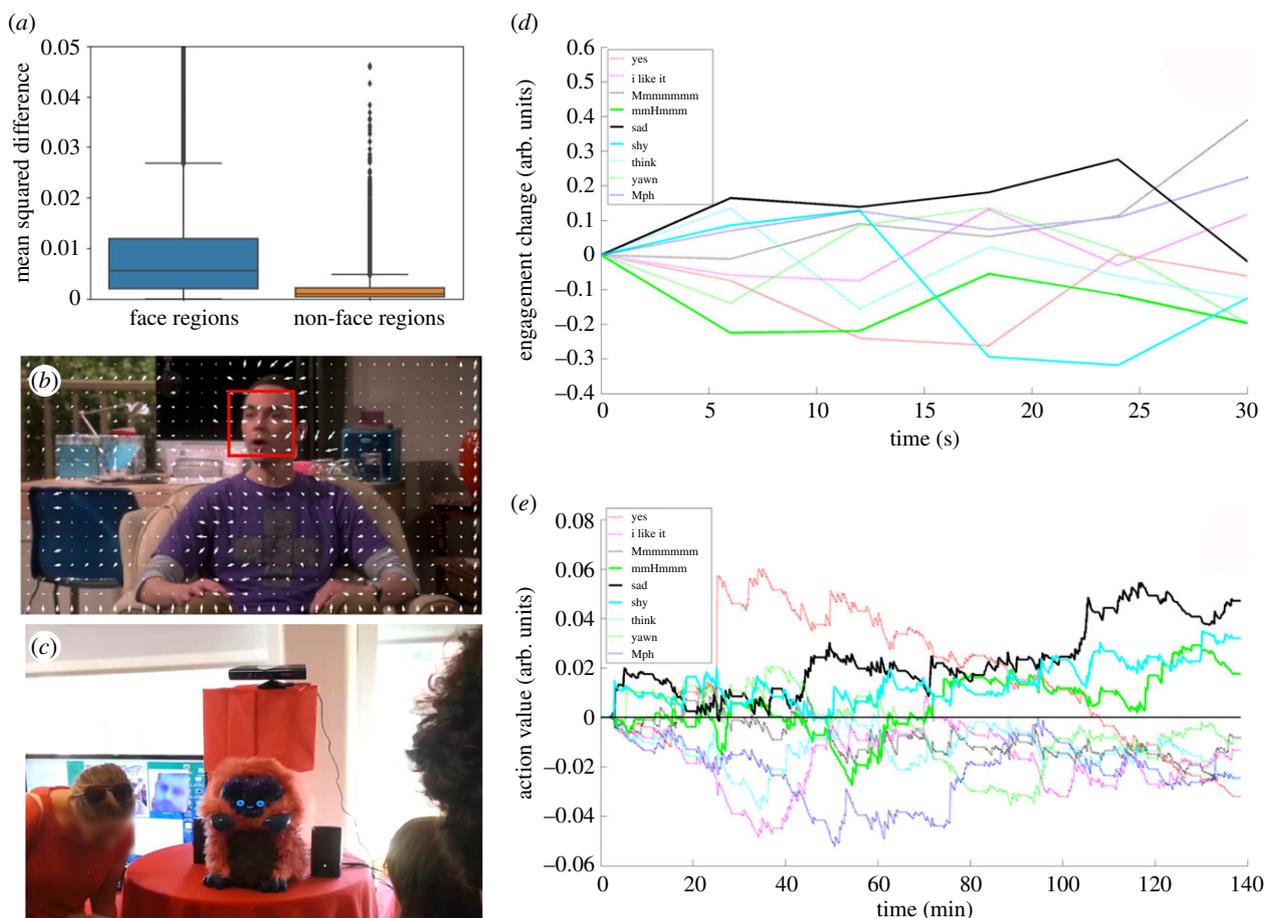


Figure 2. An embodied curious robot in social environments. (a,b) Face detection and tracking [13]. (a) The mean squared difference between two successive frames within a face region and outside of it. Cutoff at 0.05. A face region is detected using OpenCV frontal Viola-Jones face detector with 1.1 enlargement rate and with at least 6 neighbours. Performed on the first 6000 shots, only for frames where a face was detected. (b) An example image with the resultant face-detection bounding boxes and the agent's policy, depicted as arrows, shifting the focus towards the face. (c–e) Facial expression for attention [14]. (c) The embodied curious robot, Dragonbot, in the social environment. (d) Attention change as a function of time after each behaviour execution, measured as the difference in reward from the onset of the behaviour. (e) Policy-learning dynamics, measured as the action value of each action as a function of time. Dotted lines denote actions that ended with action value below zero for visualization purposes.

In essence, the Deep Curiosity Loop presented in [13] shows one possible developmental path, wherein a curious agent can learn to detect faces solely owing to their informative nature and intrinsic motivation that leads to valuing such information. This shows that being a curious agent embedded in a social environment results in the emergence of social prerequisites, such as face detection and tracking. What could a socially curious robot do if it could influence the social scenario?

(c) Facial expressions for attention

Inspired by the evidence that faces are sources of information and the emergence of face detection from the curiosity algorithm, Gordon & Breazeal [18] turned to 'socially-active' behaviours. The embodied curious agent employed was a DragonBot, which is a very expressive social robotic platform that has a large repertoire of possible facial expressions and actions, such as happy, sad, yawn, thinking, nodding, etc. [36]. It was augmented with an external camera as its sensor (figure 2c). Since curiosity entails maximizing information, the curious robot should act to receive as much information as it can. Since faces are sources of information, the robot should thus behave to have as many faces interacting with it for the longest period of time. In this constellation, this amounts to 'which facial expression will grab the most attention from people?'

The Dragonbot was controlled by a stateless reinforcement-learning algorithm, whose action-space was composed of its facial expressions and it received rewards based on the number of faces detected within the field of vision of its camera. The Dragonbot was deployed in an extremely crowded scenario, during the World Science Festival (WSF) in New York on 1 June 2014.

Examining the raw data revealed a known truth about human empathy, namely, that the 'sad' face made people stay near the robot the longest, compared with the other facial expressions (figure 2d). The curiosity algorithm picked up on this human facet after only 2 h of learning and was 'crying', i.e. presented the crying facial expression, more than all the other facial expressions (figure 2e).

5. Discussion

The developmental paths presented in this contribution showed the importance of embodied curiosity. When embedded in non-social environments, the curious robot learned about itself, i.e. its own body and sensorimotor correlations and also started to explore moving objects via reaching. By contrast, embedding within a social environment resulted in a robot that learned much more complex and intricate social behaviours, e.g. crying to get attention. The same

relatively simple algorithm, namely, rewarding information, was executed in both paths, yet the emergent behaviours differed radically. The source of the difference was the source of information: in the non-social scenario, the robot itself was the only source of information, emphasizing the importance of the action–perception cycle [37]; in the social scenario, people supplied the information, thus forcing the robot to learn more complex dynamics and behaviours.

While we have shown how social behaviours can emerge from an embodied curious agent, there is still much to be studied.

(a) Non-verbal versus verbal interaction

In this contribution, we have completely ignored verbal communication, which is a major source of social interaction. While artificial curiosity has started to be implemented in the field of verbal communication [20], there is still much to be explored. For example, the emergence of the first spoken words is still a mystery [38]. Is socially embodied curiosity its sole source?

Higher-level verbal curiosity, e.g. asking questions, is another important research venue that holds great promise in the field of education [39]. Is the order of questions asked by children dictated by the Curiosity Loop paradigm or are there other underlying mechanisms?

While curious social robots have been shown to be able to promote curiosity [40] and growth-mindset in children [41], insights into the emergence of verbal curious behaviour are still lacking.

(b) Multi-modal learning

We have explored visual developmental paths, yet multi-modal curiosity is an extremely important aspect of embodied curiosity. Visual–auditory correlation is a major path to social interaction and communication [42]. While they share important aspects with the aforementioned verbal component, non-verbal auditory perception and generation, e.g. vocal social cues, are also of great importance in social interaction [43]. How are they learned? Do they emerge as part of an information-rewarding system?

The tactile domain is another crucial developmental path in infants, and while it has been explored previously within

the artificial curiosity field [22], insights into the multi-modal integration of tactile and visual information streams are still lacking.

(c) Neuroscience

We have focused in this contribution on computational models and robotic implementation. However, major advances have been made on the underlying neuronal mechanisms of curiosity [16,44].

Areas that were associated with curiosity, during an anticipatory period of receiving desired information, were the left caudate nucleus, bilateral inferior frontal gyrus (IFG) and loci in the putamen and globus pallidus [44]. Mid-brain dopaminergic (DA) cells and cells in the orbitofrontal cortex (OFC), a pre-frontal area that receives DA innervation, were shown to encode the anticipation of obtaining reliable information from visual cues in non-human primates [45,46].

These studies suggest that the dopaminergic circuits, which have been implicated as part of the reward system and have been linked to reinforcement learning [47], are also involved in the curiosity network. However, a full account of the Curiosity Loop is yet to be examined, wherein the relationship between the actual learning process, the reward and the change in behaviour is linked. Furthermore, studies into the *developmental* aspects of curiosity are still missing in order to test the relevance of the proposed hierarchical Curiosity Loops architecture.

6. Conclusion

Curious robots hold great promise for a better understanding of human development. The bottom-up constructivist approach presented here shows that embodied curiosity in social environments may account for the emergence of social behaviours. Whether and how much other major factors influence it, e.g. genetics and predisposition, remain to be seen.

Data accessibility. This article has no additional data.

Competing interests. I declare I have no competing interests.

Funding. G.G. is a Jacobs Foundation Fellow.

References

- Piaget J. 1965 *The moral judgment of the child*. New York, NY: Free Press.
- Halford GS. 2014 *Children's understanding: the development of mental models*. New York, NY: Psychology Press.
- Fotopoulou A, Tsakiris M. 2017 Mentalizing homeostasis: the social origins of interoceptive inference. *Neuropsychanalysis* **19**, 3–28. (doi:10.1080/15294145.2017.1294031)
- Sliwa J, Freiwald WA. 2017 A dedicated network for social interaction processing in the primate brain. *Science* **356**, 745–749. (doi:10.1126/science.aam6383)
- Weng J. 2004 Developmental robotics: theory and experiments. *Int. J. Humanoid Rob.* **1**, 199–236. (doi:10.1142/S0219843604000149)
- Cangelosi A, Schlesinger M. 2015 *Developmental robotics: from babies to robots*. Cambridge, MA: MIT Press.
- Schmidhuber J. 1990 *A possibility for implementing curiosity and boredom in model-building neural controllers*. Cambridge, MA: MIT Press.
- Oudeyer PY, Kaplan F, Hafner VV. 2007 Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evol. Comput.* **11**, 265–286. (doi:10.1109/TEVC.2006.890271)
- Schmidhuber J. 2010 Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Trans. Auton. Ment. Dev.* **2**, 230–247. (doi:10.1109/TAMD.2010.2056368)
- Gordon G, Ahissar E. 2012 A curious emergence of reaching. In *Advances in autonomous robotics, TAROS 2012*. Lecture Notes in Computer Science, vol. 7429, pp. 1–12. Berlin, Germany: Springer.
- Saegusa R, Metta G, Sandini G, Natale L. 2014 Developmental perception of the self and action. *IEEE Trans. Neural Netw. Learn. Syst.* **25**, 183–202. (doi:10.1109/TNNLS.2013.2271793)
- Gordon G. 2011 Hierarchical exhaustive construction of autonomously learning networks. *Evolutionary and Reinforcement Learning for Autonomous Robot Systems, ERLARS*.
- Barkan J, Gordon G. 2018 Deep curiosity loops in social environments. (<http://arxiv.org/abs/1806.03645>)
- Gordon G, Fonio E, Ahissar E. 2014 Learning and control of exploration primitives. *J. Comput.*

- Neurosci.* **37**, 259–280. (doi:10.1007/s10827-014-0500-1)
15. Gordon G (ed.). 2018 *The new science of curiosity*. Hauppauge, NY: Nova Science Publishers Inc.
 16. Gottlieb J, Lopes M, Oudeyer P-Y. 2016 Motivated cognition: neural and computational mechanisms of curiosity, attention, and intrinsic motivation. In *Recent developments in neuroscience research on human motivation*. Advances in Motivation and Achievement, vol. 19, pp. 149–172. Bingley, UK: Emerald Group Publishing Limited.
 17. Pathak D, Agrawal P, Efros AA, Darrell T. 2017 Curiosity-driven exploration by self-supervised prediction. (<http://arxiv.org/abs/1705.05363>)
 18. Gordon G, Breazeal C. 2014 Learning to maintain engagement: no one leaves a sad Dragonbot. In *AAAI Fall Symp. Series, Artificial Intelligence for Human-Robot Interaction, The AAAI Fall Symposium, November 2014*. Technical Report FS-14-01.
 19. Gordon G, Fonio E, Ahissar E. 2014 Emergent exploration via novelty management. *J. Neurosci.* **34**, 12 646–12 661. (doi:10.1523/JNEUROSCI.1872-14.2014)
 20. Moulin-Frier C, Oudeyer P-Y. 2012 Curiosity-driven phonetic learning. In *2012 IEEE Int. Conf. on Development and Learning and Epigenetic Robotics (ICDL), San Diego, CA, 7–9 November 2012*, pp. 1–8. IEEE.
 21. Stulp F, Oudeyer PY. 2012 Emergent proximo-distal maturation through adaptive exploration. In *2012 IEEE Int. Conf. on Development and Learning and Epigenetic Robotics (ICDL), San Diego, CA, 7–9 November 2012*, pp. 1–6. IEEE.
 22. Pape L, Oddo CM, Controzzi M, Cipriani C, Förster A, Carrozza MC, Schmidhuber J. 2012 Learning tactile skills through curious exploration. *Front. Neurobot.* **6**, 6. (doi:10.3389/fnbot.2012.00006)
 23. Lalazar H, Vaadia E. 2008 Neural basis of sensorimotor learning: modifying internal models. *Curr. Opin. Neurobiol.* **18**, 573–581. (doi:10.1016/j.conb.2008.11.003)
 24. Wolpert DM, Kawato M. 1998 Multiple paired forward and inverse models for motor control. *Neural Netw.* **11**, 1317–1329. (doi:10.1016/S0893-6080(98)00066-5)
 25. Innocenti GM, Price DJ. 2005 Exuberance in the development of cortical networks. *Nat. Rev. Neurosci.* **6**, 955–965. (doi:10.1038/nrn1790)
 26. Pinto L, Gupta A. 2016 Supersizing self-supervision: learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE Int. Conf. on Robotics and Automation (ICRA), Stockholm, Sweden, 16–21 May 2016*, pp. 3406–3413. IEEE.
 27. Ivaldi S, Nguyen SM, Lyubova N, Droniou A, Padois V, Filliat D, Oudeyer P, Sigaud O. 2014 Object learning through active exploration. *IEEE Trans. Auton. Ment. Dev.* **6**, 56–72. (doi:10.1109/TAMD.2013.2280614)
 28. McDuff D, Kaliouby R, Senechal T, Amr M, Cohn J, Picard R. 2013 Affectiva-MIT facial expression dataset (AM-FED): naturalistic and spontaneous facial expressions collected. *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition Workshops 2013*, pp. 881–888. IEEE.
 29. Bagautdinov TM, Alahi A, Fleuret F, Fua P, Savarese S. 2017 Social scene understanding: end-to-end multi-person action localization and collective activity recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 22–25 July 2017, Honolulu, HI*, pp. 3425–3434. IEEE.
 30. Jayaraman S, Fausey CM, Smith LB. 2015 The faces in infant-perspective scenes change over the first year of life. *PLoS ONE* **10**, e0123780. (doi:10.1371/journal.pone.0123780)
 31. Fausey CM, Jayaraman S, Smith LB. 2016 From faces to hands: changing visual input in the first two years. *Cognition* **152**, 101–107. (doi:10.1016/j.cognition.2016.03.005)
 32. Viola P, Jones MJ. 2004 Robust real-time face detection. *Int. J. Comput. Vision* **57**, 137–154. (doi:10.1023/B:VISI.0000013087.49260.fb)
 33. Jakobsen KV, Umstead L, Simpson EA. 2016 Efficient human face detection in infancy. *Dev. Psychobiol.* **58**, 129–136. (doi:10.1002/dev.21338)
 34. van Hasselt H, Guez A, Silver D. 2015 Deep reinforcement learning with double Q-learning. (<http://arxiv.org/abs/1509.06461>).
 35. Cichy RM, Khosla A, Pantazis D, Torralba A, Oliva A. 2016 Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci. Rep.* **6**, 27755. (doi:10.1038/srep27755)
 36. Setapen AM. 2012 Creating robotic characters for long-term interaction. PhD dissertation, Massachusetts Institute of Technology.
 37. Gordon G, Kaplan DM, Lankow B, Little DY, Sherwin J, Suter BA, Thaler L. 2011 Toward an integrated approach to perception and action: conference report and future directions. *Front. Syst. Neurosci.* **5**, 20. (doi:10.3389/fnsys.2011.00020)
 38. Roy BC, Frank MC, DeCamp P, Miller M, Roy D. 2015 Predicting the birth of a spoken word. *Proc. Natl Acad. Sci. USA* **112**, 12 663–12 668. (doi:10.1073/pnas.1419773112)
 39. Jirout JJ. 2011 Curiosity and the development of question generation skills. In *2011 AAAI Fall Symp. Series, Question Generation, Arlington, VA, 4–6 November 2011*, Technical Report FS-11-04. AAAI.
 40. Gordon G, Breazeal C, Engel S. 2015 Can children catch curiosity from a social robot? In *Proc. of the Tenth Annual ACM/IEEE Int. Conf. on Human-Robot Interaction, Portland, OR, 2–5 March 2015*, pp. 91–98. New York, NY: ACM.
 41. Park HW, Rosenberg-Kima R, Rosenberg M, Gordon G, Breazeal C. 2017 Growing growth mindset with a social robot peer. In *Proc. of the 2017 ACM/IEEE Int. Conf. on Human-Robot Interaction, HRI '17, Vienna, Austria, 6–9 March 2017*, pp. 137–145. New York, NY: ACM.
 42. Cichy RM, Teng S. 2017 Resolving the neural dynamics of visual and auditory scene processing in the human brain: a methodological approach. *Phil. Trans. R. Soc. B* **372**, 20160108. (doi:10.1098/rstb.2016.0108)
 43. Manusov V, Trees AR. 2002 Are you kidding me: the role of nonverbal cues in the verbal accounting process. *J. Commun.* **52**, 640–656. (doi:10.1111/j.1460-2466.2002.tb02566.x)
 44. Kang MJ. 2010 Three experimental studies of reward and decision making. Doctoral dissertation, California Institute of Technology.
 45. Blanchard TC, Hayden BY, Bromberg-Martin ES. 2015 Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* **85**, 602–614. (doi:10.1016/j.neuron.2014.12.050)
 46. Ethan S. 2009 Bromberg-Martin and Okihide Hikosaka. Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* **63**, 119–126. (doi:10.1016/j.neuron.2009.06.009)
 47. Montague PR, Dolan RJ, Friston KJ, Dayan P. 2012 Computational psychiatry. *Trends. Cogn. Sci.* **16**, 72–80. (doi:10.1016/j.tics.2011.11.018)