

Designing Toys That Come Alive: Curious Robots for Creative Play

Kathryn Merrick

School of Information Technologies and Electrical Engineering
University of New South Wales, Australian Defence Force Academy
Northcott Drive, Canberra ACT, 2612,
k.merrick@adfa.edu.au

Abstract. Creative thinking requires imagination, creativity, play, sharing and reflection. This paper presents an architecture for a curious, reconfigurable robot that encourages creative design thinking by permitting designed structures to learn behaviours. These behaviours encourage designers to play with different structures, reflect on the relationship between structure and behaviour and imagine new structures. A demonstration of the architecture is described using the Lego Mindstorms platform. The demonstration shows how a curious robot can adapt new behaviours in response to changes in its structure, and how this can encourage the creative thinking spiral and creative design.

Keywords: Curiosity, motivated reinforcement learning, robots, creative design, creative play.

1 Creativity and Design

Design can be described as a process of purposeful, constrained decision making requiring exploration and learning [2]. The role of creativity in design has many interpretations, including the distinction between the design of creative artefacts and the evaluation of the processes involved in design as creative [1, 3, 10]. Renick [7] writes that there has been a transition from the Industrial Society to the Information Society and, most recently, the Creative Society. He identifies the need to teach people to use creative thinking processes involving imagination, creativity, play, sharing and reflection.

This paper presents an architecture for curious, reconfigurable robots as a toy that encourages creative play. The architecture permits designed structures to learn behaviours to encourage play, reflection and imagination for ongoing, creative design.

The remainder of this section discusses reconfigurable robots as toys and introduces the idea of curious robots to encourage creative play. Section 2 presents an architecture for curious, reconfigurable robots using the Lego Mindstorms platform. Section 3 describes an initial demonstration of a curious robot learning behaviour in response to changes in its structure. The paper concludes by reflecting on how curious robots can encourage creative design by influencing the creative thinking spiral [7].

1.1 Reconfigurable Robots and Creative Play

Reconfigurable robots comprise sets of modules that can be re-arranged to achieve different structures, behaviours and functions. Reconfigurable robots have been developed for their functional, economic and creative advantages. Functionally, reconfigurable robots promise engineering versatility, flexibility, robustness and the ability to self-repair through redundancy [13]. Likewise, economic advantage can be gained by designing complex machines as reconfigurable sets of modules.

From a creative viewpoint, technologies for reconfigurable robots include toys such as Lego (<http://www.lego.com>) and Meccano (<http://www.meccano.com>). These products encourage creative play, creative thinking and creative design by providing sets of basic components and connector modules that can be built into different structures. Recent versions of these toys can be fitted with motorised and even electronic components to control the behaviour of the final structure. Meccano creations can be fitted with motors or remote controls and Lego Mindstorms products can be attributed behaviour using a programmable brick.

The incorporation of motors, remote controls and programmable bricks opens the way traditional building packages such as Lego and Meccano to become platforms for playful learning of concepts in electronics and computer programming. In contrast, this paper proposes an architecture that draws on the programmable capacity of reconfigurable robots to encourage creative design thinking through creative play. This architecture uses computational models of curiosity and machine learning to create curious robots that are able to ‘come alive’ with new behaviours in response to changes in their structure. The aim is to encourage experimentation with structure by providing the creative designer with real-time information about the relationship between the structure and its possible behaviours.

1.2 Curious Robots

While existing research has focused on developing hardware modules for reconfigurable robots, and the software required for those modules to communicate, the problem of attributing behaviour to reconfigurable robots remains a challenge. Traditional artificial intelligence techniques tend to assume a fixed set of inputs and outputs (sensors and actuators) and a fixed set of goals based on these inputs and outputs [8]. This is in contrast to the needs of reconfigurable robots, which may have changing sensors and actuators and thus changing goals.

Recent work has focused on the design of developmental learning algorithms that can generate goals online in response to the changing experiences of a robot [6, 9]. These approaches use computational models of intelligent, adaptive, curiosity to develop new goals, but do not provide a way for robots to adapt to changes in their sensors or actuators. A later approach that includes adaptable representations for input and output was proposed by Merrick and Maher [5]. This paper adapts the Merrick and Maher [5] model for motivated reinforcement learning agents to design an architecture for curious, reconfigurable robots that can develop new behaviours as a response to their changing structure.

2 Architecture for a Curious, Reconfigurable Robot

The architecture presented in this section comprises two layers, a device layer and an agent layer, as shown in Figure 1. This paper uses the Lego Mindstorms NXT robot platform running the Lejos (<http://lejos.sourceforge.net>) Java firmware. However the concept of a reconfigurable robot using curiosity and motivated reinforcement learning is general enough to be applied to other reconfigurable robot hardware.

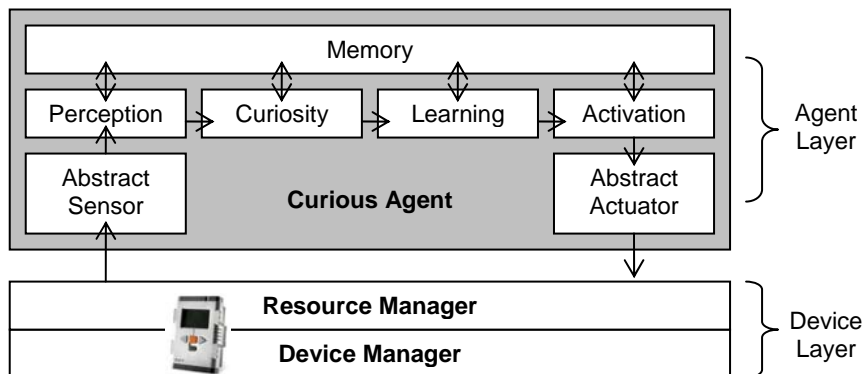


Fig. 1. Architecture for a curious, reconfigurable robot using the Lego Mindstorms platform.

2.1 Device Layer

Lego Mindstorms robots can be designed with a range of different sensors and actuators including sensors for colour, light, sound or distance to objects and actuators for servo motors. NXT sensors and actuators are heterogeneous and return different numbers and types of outputs. The device layer constructs a standardised, context-free grammar (CFG) [5] representation of sensor data and available actions and communicates this to the agent layer.

Device Manager: The device manager identifies the sensor devices $S_1, S_2, S_3 \dots$ attached to the robot. Data from all sensors is encapsulated as a single sensation S represented as a variable length string of label-value pairs. The device manager also identifies the actuators $A_1, A_2, A_3 \dots$ attached to the robot.

Resource Manager: The resource manager communicates sensations and actions between the agent and the devices via Bluetooth. In the demonstration in Section 3, the agent layer is run on a PC separate from the NXT brick. The Lejos Java firmware does not currently offer sufficient support for the complex programming constructs required to implement a curious agent. In future it is envisaged that it will be possible to run the agent layer on the NXT brick and the resource manager will no longer be required.

2.2 Agent Layer

The agent layer comprises a motivated reinforcement learning agent [5] using a computational model of curiosity as the motivation function. The agent has four processes for perception, curiosity, learning and activation. These communicate with the physical sensors and actuators via the device layer.

Abstract Sensor: The abstract sensor communicates with the resource manager via Bluetooth and receives a sensation S and set A of available actions at each time-step t .

Perception: The perception process computes an event representing the change between the current sensation and the previous sensation stored in memory. Events are also represented as variable length strings of label-value pairs. The value of each event element is computed as the real number difference of sensation elements with the same label.

Curiosity: The curiosity process computes a curiosity value for the current event. Curiosity is computed as a function of the novelty of an event. Novelty is calculated using an Habituated Self-Organising Map [4] such that novelty depends on the similarity of the current event to previously experienced events. The curiosity value is computed using the Wundt [12] curve, shown in Figure 2, so that the most curious events are those that are moderately novel in the agent's experience.

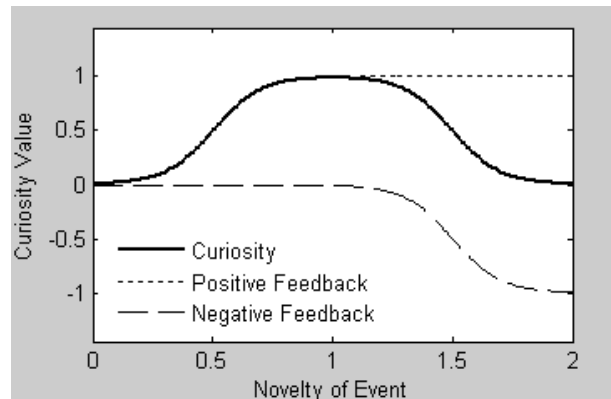


Fig. 2. Model of curiosity using the Wundt [12] curve.

Learning: The learning process uses Q-learning [10] to update a state-action table stored in memory. This table maps sensations to actions and utility values. Utility is computed using the curiosity value as the reward in the Q-learning update. Because the agent will compute high curiosity for different events at different times based on its experiences with its current sensors and actuators, the agent will learn different behaviours that are adapted to its current sensors and actuators and current structure.

Activation: The activation process uses an ϵ -greedy algorithm to select the action with the highest utility value from the state-action table 90% of the time and a random action 10% of the time. This means that the agent follows its learned behaviours 90% of the time and experiments 10% of the time. This encourages the robot to adapt.

Abstract Actuator: The abstract actuator communicates the unique identifier of the action selected by the activation process to the resource manager.

3 Demonstration of a Curious Robot for Creative Play

This section describes a proof-of-concept demonstration of a curious, reconfigurable robot. A designer plays with the robot and changes its structure. The robot adapts its behaviour as different sensors and actuators are added, inspiring the designer to further modify the structure. The designer builds a jointed limb as shown in Figure 3. Figure 2(a) shows the lower part of a limb, mounted on a turn-table with an attached compass sensor. In this structure the curious robot learns behaviours to rotate the limb, as the agent layer is curious about changes in the compass reading.

The designer then develops the limb further into a jointed arm, as shown in Figure 2(b). The arm can extend and retract, has a touch sensor as a finger and an accelerometer to measure its tilt. In this structure, the robot is curious about changes in the accelerometer reading and learns behaviours to extend and retract the arm.

The designer then restructures the limb into a jointed neck, as shown in Figure 2(c). The neck can extend and retract and has an ultrasonic distance sensor. In this structure, the robot is curious about extension and retraction of the neck, which cause changes in the distance readings. The robot also develops a somewhat creative behaviour in which it rapidly extends and retracts the neck. The resultant jerking of the structure causes changes in the compass sensor reading, which the agent finds curious.



Fig. 3. Reconfiguring a curious robot limb. (a) Lower limb and compass sensor (b) Jointed arm with touch sensor and accelerometer (c) Jointed neck with ultrasonic distance sensor.

The examples above show how the curious agent learn behaviours for different structures, without requiring changes to the agent model. This encourages the designer to play with relationship between structure and behaviour. However, a number of issues arising from the combination of the curious agent with the reconfigurable robot were also identified in this pilot study. First, high resolution sensors can make learning of behaviours quite slow and introduce noise that inhibits learning. In this study, for example, the resolution of the compass sensor was reduced from 360 degrees to 8 compass points to speed learning.

The second issue identified is that the physical structure of the robot may include limitations that are not programmed for in the generic curious agent. For example, the extension and retraction of the limb in this study are bounded by the maximum and minimum angle that can be achieved by the configuration of vertical and horizontal beams. However the servo motors are relatively strong, so the curious agent can experiment with the limb to destruction by forcing it past its limits. This suggests that

future reconfigurable robots will have to utilise more flexible joint structures. For example, cog setups where continuous forward (or backward) motion of a motor results in a cycle of retraction and extension.

4 Conclusion

This paper has presented an architecture for curious, reconfigurable robots for creative play. Curious robots can learn new behaviour in response to changes in their structure. The robots in this paper use a computational model of curiosity to identify new goals and reinforcement learning to develop new behaviour in response to those goals. An initial demonstration of curious robots using the Lego Mindstorms platform shows how curious, reconfigurable robots can influence the creative thinking spiral. The adaptive behaviour of the robot in response to its changing structure encourages reflection, imagination and ongoing creative play.

In future, this research may proceed in two directions. First, user studies can provide insight into the impact of this technology on creative play. In conjunction with this, the technology itself can be further developed with the aim of designing robots that can develop complex, realistic and meaningful behaviours. Such robots might have applications beyond entertainment robots to explorer robots in spaces or assistant robots in the home or industry.

References

1. Boden, M: The creative mind, myths and mechanisms, Wiedenfield and Nicholson, London, (1991)
2. Gero, J.S: Creativity, emergence and evolution in design, Second International Roundtable Conference on Computational Models of Creative Design, Australia, pp 1-28, (1992).
3. Kim, S. H: Essence of creativity, Oxford University Press, New York, (1990)
4. Marsland, S., Nehmzow, U, Shapiro, J.: A real-time novelty detector for a mobile robot, EUREL European Advanced Robotics Systems Masterclass and Conference, (2000).
5. Merrick, K., Maher, M-L.: Motivated reinforcement learning for adaptive characters in open-ended simulation games, ACM SIGCHI International Conference on Advances in Computer Entertainment Technology, (ACE 2007), Salzburg, Austria, pp 127-134.
6. Oudeyer, P-Y.: Intelligent adaptive curiosity: a source of self-development, In proceedings of the fourth international workshop on epigenetic robotics, pp 127-130, (2004)
7. Resnick, M: Sowing the seeds for a more creative society, Learning and Leading with Technology, International Society for Technology in Education, pp 18-22 (2007)
8. Russel, S., Norvig, P.: Artificial intelligence a modern approach, Prentice Hall, (1995).
9. Schmidhuber, J.: Curious model building control systems, International Joint Conference on Artificial Neural Networks. IEEE, Singapore, pp1458-1463.
10. Sternberg, R: The nature of creativity, Cambridge University Press, (1988).
11. Watkins, C., Dayan, P.: Q-learning, Machine Learning, 8(3) pp 270-292, (1992).
12. Wundt, W.: Principles of physiological psychology. Macmillan, New York, (1910).
13. Yim, M., Duff, D., Roufas, K.: Polybot: a modular, reconfigurable robot, Proceedings of the 2000 IEEE International Conference on Robotics and Automation, CA, pp 514-520 (2000)