



Editorial: Intrinsically Motivated Open-Ended Learning in Autonomous Robots

Vieri Giuliano Santucci^{1*}, Pierre-Yves Oudeyer², Andrew Barto³ and Gianluca Baldassarre¹

¹ Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche, Rome, Italy, ² Institut National de Recherche en Informatique et en Automatique (INRIA), Bordeaux, France, ³ College of Information and Computer Sciences, University of Massachusetts, Amherst, MA, United States

Keywords: intrinsic motivation, open-ended learning, robotics, developmental robotics, curiosity driven learning

Editorial on the Research Topic

Intrinsically Motivated Open-Ended Learning in Autonomous Robots

Notwithstanding the important advances in Artificial Intelligence (AI) and robotics, artificial agents still lack the necessary autonomy and versatility to properly interact with realistic environments. This requires agents to face situations that are unknown at design time, to autonomously discover multiple goals/tasks, and to be endowed with learning processes able to solve multiple tasks incrementally and online.

Starting in developmental robotics (Lungarella et al., 2003; Cangelosi and Schlesinger, 2015), and gradually expanding into other fields, intrinsically motivated learning (sometimes called “curiosity-driven learning”) has been studied by many researchers as an approach to autonomous lifelong learning in machines (Oudeyer et al., 2007; Schmidhuber, 2010; Barto, 2013; Mirolli and Baldassarre, 2013). Inspired by the ability of humans and other mammals to discover how to produce “interesting” effects in the environment driven by self-generated motivational signals not related to specific tasks or instructions (White, 1959; Berlyne, 1960; Deci and Ryan, 1985), the research in the field of intrinsically motivated open-ended learning aims to develop agents that autonomously generate motivational signals (Merrick, 2010) to acquire repertoires of diverse skills that are likely to become useful later when specific “extrinsic” tasks need to be performed (e.g., Barto et al., 2004; Baldassarre, 2011; Baranes and Oudeyer, 2013; Kulkarni et al., 2016; Santucci et al., 2016).

This Research Topic aims to present state-of-the-art research on open-ended learning in autonomous robots, with a particular focus on systems driven by intrinsic motivations (but not limited to these systems), and augments the information presented at the *Third International Workshop on Intrinsically Motivated Open-ended Learning – IMOL2017*, held in Rome, Italy, 4–6 October 2017. Although the development of autonomous artificial agents is pursued via different kinds of approaches, such as information theory (Klyubin et al., 2008; Martius et al., 2013), epigenetic robotics (Lones et al., 2016), machine learning (Machado et al., 2017), and evolutionary computation (Lehman and Stanley, 2011), intrinsically motivated open-ended learning is today a mature field producing promising research.

The field nevertheless presents many open challenges, the main ones we mention here. A first open issue of central importance for open-ended learning is how an agent should autonomously generate goals and learn policies for achieving them, so that the policies are useful for solving many new tasks that are unknown when the policies are learned. Another open challenge is to design systems that use intrinsic motivations to support learning compact representations of environment states, and hence of goals; in particular, learning compact representations that are relevant for action. Another challenge involves the continuous/discrete representation of goals. It seems plausible that low-level goals (e.g., related to postures that a robot might assume) are encoded in a

OPEN ACCESS

Edited by:

Florian Röhrbein,
Technical University of Munich,
Germany

Reviewed by:

Lola Cañamero,
University of Hertfordshire,
United Kingdom

*Correspondence:

Vieri Giuliano Santucci
vieri.santucci@istc@cnr.it

Received: 25 September 2019

Accepted: 19 December 2019

Published: 17 January 2020

Citation:

Santucci VG, Oudeyer P-Y, Barto A and Baldassarre G (2020) Editorial: Intrinsically Motivated Open-Ended Learning in Autonomous Robots. *Front. Neurobot.* 13:115. doi: 10.3389/fnbot.2019.00115

continuous space, whereas high-level goals (e.g., touch, push, hit, reach, grasp) are encoded in a discrete fashion so as to be easily composed to achieve more complex goals. If this is the case, what is the relation between these different types of goal representations? A related problem is how to design architectures that can suitably manage generating goals and learning the related skills, and storing them at different levels of granularity. Then there is the open problem of how to best re-use goals and acquired skills to accomplish novel tasks (exploitation), including how to use previous learning in order to efficiently learn new goals and skills (“transfer learning”), and how to form “chains” of interrelated, hierarchical skills (“curriculum learning”), on the basis of intrinsically motivated processes. Learning different tasks in an open-ended fashion tends to cause catastrophic interference, which is another problem that needs to be addressed. Other important challenges are related to the interaction between intrinsic motivations and other forms of “natural learning” such as social interaction: how this interaction might be connected to the development of higher-level cognitive skills, e.g., language.

The contributions collected for this Topic, reviewed below, not only extend the core research in the field but also tackle some of the open questions mentioned above, with a particular focus on: the autonomous acquisition of skills and motor behaviors; the analysis of architectures and learning signals needed to perform sequences of different tasks; the interaction of artificial agents with their environments through different sensors and actuators; the formation and representation of goals; and the interplay between intrinsic motivations and other learning strategies such as imitation learning. Following is a summary of how the contributions in this collection address these challenges.

In Rayyes et al. the authors propose a system enabling a robot to learn to reach to different points in space by exploiting symmetry properties of the actuators to allow exploration to be limited to only a small part of the configuration space. Maestre et al. tackle skill learning at the level of object manipulation, where low-level and high-level features are extracted through task-agnostic interactions with the environment directed toward learning affordances that, in turn, guide the robot to solve assigned tasks. Baldassarre et al. propose a system that uses intrinsic motivations to learn forward models and affordances, here intended as the probabilities of achieving the goals of the affordance-related actions. In particular, this work examines how active-vision, which allows factoring the environment state into pieces of information related to single objects, can support such learning processes and also facilitate solving extrinsic tasks involving multiple objects through one-step planning. Learning progress might also be used as in the contribution of Uchibe to train multiple skills in parallel. In particular, here transfer learning techniques are combined with “mixture of experts” strategies to develop multiple control modules that are then used to solve control tasks with simulated agents. When learning many different tasks, an agent might try to optimize all of them simultaneously. Abdelfattah et al. leverage intrinsic motivations to develop a method that can cope with multi-objective Markov decision processes, and they compare it to other state-of-the-art algorithms. In real environments, complex

tasks might need to be learned through exposure to sequences of simpler tasks. This is a crucial issue for autonomous robotics, which is tackled in the work of Duminy et al. using an active learning approach. These authors propose a new algorithm that allows a robot to autonomously discover how to combine pre-defined primitive motor policies to learn increasingly complex combinations of motor policies. In particular, while it is learning, and agent is able to decide which outcome on which to focus and which exploration strategy to apply, leveraging imitation learning, goal babbling, and strategic learning techniques based on intrinsic motivations.

When an agent can autonomously generate different kinds of goals, task-specific reward functions might be complicated to design since they require significant domain knowledge. Intrinsic motivations might provide general, task-agnostic reward functions able to exploit the inherent properties of different goals. Moreover, as shown in Dhakan et al., these reward functions can be used as building blocks to generate sequences of tasks enabling more complex behaviors to be learned. A well-known problem, mentioned above, related to learning multiple tasks is that of catastrophic forgetting. This problem might be even harder to tackle for artificial agents that have to perform life-long learning in complex and unknown environments. Instead of constraining the set of inputs at design time, in Parisi et al., the authors propose a dual-memory self-organizing architecture with two growing recurrent networks that in parallel learn episodic memory and semantic memory, expanding their structures in response to novel sensory experiences.

Vision plays an important role in many aspects related to autonomous learning and exploration. de La Bourdonnaye et al. followed a developmental perspective and built an artificial system that learns to reach for objects in different locations in the environment by leveraging a weakly-supervised stage-wise procedure. Learning to reach is divided into three tasks: learn to fixate objects, learn hand-eye coordination (learn to fixate on the end-effector), and learn to use the previously acquired knowledge to perform reaching to different locations. Visuo-motor coordination is also tackled by Wijesinghe et al., where predictive models are used to guide a humanoid robot in learning to track its hand and other movements in the visual field without the use of any forward kinematics or pre-defined visual feature descriptors. The use of prediction in multi-sensory integration allows a better incorporation of proprioceptive and visual cues and leads to the development of emergent properties similar to those of human hand-eye coordination. Task-agnostic motivations such as information gain are used in the work of Dauce to drive action selection and the exploration of visual inputs: promising results are shown, highlighting how compression strategies might improve both performance in visual recognition and efficiency of the system thanks to reduced computational costs. Autonomous exploration is the focus of Cohen-Lhyver et al. These authors underline the important role of attention, which they claim can be considered as a sort of intrinsic motivation. They implemented this notion in a humanoid robot and showed how two components (congruence and reduction of uncertainty) can be used to explore new environments following audio-visual inputs encoded at a semantic level.

Similar to appealing to information gain, exploration drive by “criticality” can be used to generate autonomous behaviors. Aguilera and Bedia explore this connection through conceptual models that exploit maximum entropy to drive agents toward critical points (e.g., transition points between different kinds of behaviors). Finally, Mahzoon et al. focus on the development of social robots. Within a developmental perspective, these authors address problems related to training real-world robots by presenting two new algorithms that improve a robot’s performance in terms of learning efficiency, complexity of the learned behaviors, and predictability of the robot’s behavior.

In addition to the aforementioned 14 original research articles, four additional papers have been published within this Research Topic. The first is a methodological article in which Yu et al. present an algorithm that takes into account the constructive interplay between boredom and curiosity, giving rise to effective exploration and forward model learning. The second is a review article in which Khan et al. examine the motivational systems used in computational models to build agents capable of autonomous goal generation and task learning. The authors then investigate how these strategies might be transferred to multi-agent systems and swarms, highlighting the current state-of-the-art and future key challenges. The last two papers are perspective articles. In the first of these, Doncieux et al. argue that a key issue for an agent performing open-ended

learning is not only the problem of maximizing the rewards related to the different tasks, but also the problem of building proper representations of the states and the actions describing the tasks themselves. The authors present a conceptual framework to address this crucial issue, underlining the central role of intrinsic motivations. In the second article, Palm and Schwenker analyse the use of reinforcement learning (RL) in the field of developmental robotics, describing its strengths and weaknesses with respect to some specific problems that arise in the field. The authors suggest that multi-objective RL might face some of the problems they listed and that leveraging multiple motivations can improve RL agents’ learning performance.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

FUNDING

The research and editorial activity related to this topic, has been partially funded from the European Union’s Horizon 2020 Research and Innovation Program under Grant Agreement no. 713010 (GOAL-Robots—Goal-based Open-ended Autonomous Learning Robots).

REFERENCES

- Baldassarre, G. (2011). “What are intrinsic motivations? a biological perspective,” in *Proceedings of the International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob-2011)* (New York, NY; Frankfurt am Main: IEEE), E1–E8.
- Baranes, A., and Oudeyer, P.-Y. (2013). Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robot. Auton. Syst.* 61, 49–73. doi: 10.1016/j.robot.2012.05.008
- Barto, A. G. (2013). “Intrinsic motivation and reinforcement learning,” in *Intrinsically Motivated Learning in Natural and Artificial Systems* (Berlin; Heidelberg: Springer), 17–47.
- Barto, A. G., Singh, S., and Chentanez, N. (2004). “Intrinsically motivated learning of hierarchical collections of skills,” in *Proceedings of the 3rd International Conference on Development and Learning* (La Jolla, CA), 112–119.
- Berlyne, D. E. (1960). *Conflict, Arousal, and Curiosity*. New York, NY: McGraw Hill.
- Cangelosi, A., and Schlesinger, M. (2015). *Developmental Robotics: From Babies to Robots*. Cambridge, MA: MIT Press.
- Deci, E. L., and Ryan, R. M. (1985). *Intrinsic Motivation and Self-Determination in Human Behavior*. New York, NY: Plenum Press.
- Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2008). Keep your options open: an information-based driving principle for sensorimotor systems. *PLoS ONE* 3:e4018. doi: 10.1371/journal.pone.0004018
- Kulkarni, T. D., Narasimhan, K., Saeedi, A., and Tenenbaum, J. (2016). “Hierarchical deep reinforcement learning: integrating temporal abstraction and intrinsic motivation,” in *Advances in Neural Information Processing Systems* (Bacelona), 3675–3683.
- Lehman, J., and Stanley, K. O. (2011). Abandoning objectives: evolution through the search for novelty alone. *Evol. Comput.* 19, 189–223. doi: 10.1162/EVCO_a_00025
- Lones, J., Lewis, M., and Cañamero, L. (2016). From sensorimotor experiences to cognitive development: investigating the influence of experiential diversity on the development of an epigenetic robot. *Front. Robot. AI* 3:44. doi: 10.3389/frobt.2016.00044
- Lungarella, M., Metta, G., Pfeifer, R., and Sandini, G. (2003). Developmental robotics: a survey. *Connect. Sci.* 15, 151–190. doi: 10.1080/09540090310001655110
- Machado, M. C., Bellemare, M. G., and Bowling, M. (2017). “A laplacian framework for option discovery in reinforcement learning,” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2295–2304.
- Martius, G., Der, R., and Ay, N. (2013). Information driven self-organization of complex robotic behaviors. *PLoS ONE* 8:e63400. doi: 10.1371/journal.pone.0063400
- Merrick, K. E. (2010). A comparative study of value systems for self-motivated exploration and learning by robots. *IEEE Trans. Auton. Ment. Dev.* 2, 119–131. doi: 10.1109/TAMD.2010.2051435
- Mirolli, M., and Baldassarre, G. (2013). “Functions and mechanisms of intrinsic motivations,” in *Intrinsically Motivated Learning in Natural and Artificial Systems*, eds G. Baldassarre and M. Mirolli (Berlin; Heidelberg: Springer), 49–72.
- Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Trans. Evol. Comput.* 11, 265–286. doi: 10.1109/TEVC.2006.890271
- Santucci, V. G., Baldassarre, G., and Mirolli, M. (2016). Grail: a goal-discovering robotic architecture for intrinsically-motivated learning. *IEEE Trans. Cogn. Dev. Syst.* 8, 214–231. doi: 10.1109/TCDS.2016.2538961
- 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
- White, R. W. (1959). Motivation reconsidered: the concept of competence. *Psychol. Rev.* 66:297. doi: 10.1037/h0040934

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Santucci, Oudeyer, Barto and Baldassarre. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.