Guest Editorial
Special Issue on Intrinsically Motivated Open-Ended Learning (IMOL)

I. Scope of This Special Issue

THE CONCEPT of lifelong [1], continual [2], progressive [3], or open-ended [4], [5], [6] learning by artificial agents or robots is of interest to researchers because it permits robots to adapt to multiple tasks over the course of their life and progressively accumulate knowledge [1], [3], [7]. This means that the agent or robot is less likely to become obsolete due to environmental changes, and it may be better equipped to respond to new tasks through the accumulated knowledge. This can include motor skills [8], high-level behaviors [9], vision [10], and social skills [11], among others. On the other hand, the ability to model and thus "understand" lifelong learning is of potential relevance for robots that interact with children or the elderly, or software agents that interact with people in education or training settings.

The study of autonomous, open-ended learning in artificial agents started as a research niche in the fields of machine learning and developmental robotics. The field uses inspiration from biological, cognitive, and social systems research on topics, such as intrinsic motivations, control striving, and goal-directed behavior to develop systems able to autonomously explore an environment and learn skills driven by self-generated goals.

With the improvements in machine learning techniques, such as the emergence of deep learning, artificial intelligence and machine learning are finding more and more applications in every day life, commerce and industry in both software and robotics systems. Despite this, the pursuit of autonomous artificial agents, and the applications in which it is acceptable for them to operate, remains an open challenge. Artificial agents are still limited in their flexibility and ability to adapt to a wide range of circumstances. This prevents them from interacting with realistic environments where they have to face situations that were not predicted by the system designer; where learning needs to address multiple tasks, or when those tasks have to be discovered and solved online and autonomously.

This special issue brings together the most recent state-of-the-art research in intrinsically motivated, open-ended learning for autonomous agents and developmental robots.

II. Contributions to the Special Issue

This special issue includes nine papers that cover a breadth of topics in open-ended learning. The first three papers address architectures and algorithms for intrinsically motivated learning; the next paper investigates the combination of extrinsic and intrinsic motives to improve task-oriented learning performance; a further two papers address social aspects of motivation incorporating human teachers and the final three papers look at applications, including dialog development, object recognition, and tactile exploration.

The first three papers address architectures and algorithms for intrinsically motivated learning. The first by Sener et al. [A1] focuses on self-exploration. An exploration mechanism is proposed that in which a latent space combines actions, objects, and action–outcome representations. Local regions are formed in the late space to permit forward model learning. At a given learning step, an agent uses intrinsic motivation to select the forward model with the highest learning progress to adapt. This approach is inspired by the way infants learn, because high learning progress in a particular region implies that the learning problem is neither too easy nor too difficult. The proposed approach is validated on a simulated robot interacting with objects on a tabletop. The robot learns the outcomes of these interactions and organizes its own learning curriculum. The authors discuss how the learning mechanism demonstrates features that are aligned with certain aspects of infant development. In particular, the proposed system learns to predict grasp action outcomes earlier than that of push action.

These second of these papers by Zhang et al. [A2] focuses on a model of intrinsic plasticity. Intrinsic plasticity, first identified in biological nerve cells, is an unsupervised, self-adaptive, local learning rule. It has been shown to be able to maximize the transmission entropy of neuronal information. This paper proposes the soft-reset leaky integrate-and-fire model with a new intrinsic plasticity learning rule that optimizes the neuronal membrane potential state to be exponentially distributed. The proposed soft-reset model can avoid the problem that conventional leaky integrate-and-fire neuronal membrane potential is not fully differentiable, hence the proposed intrinsic plasticity rule can directly regulate the membrane potential as an auxiliary output signal.

The third contribution to architecture and algorithm design revolves around abstract representation learning by Wilmot et al. [A3]. A key need for open-ended learning is the ability to form increasingly abstract representations. This drives the development of complex behavior. This paper considers the learning of abstract representations in a multimodal setting. The problem is treated as a lossy compression problem

Digital Object Identifier 10.1109/TCDS.2023.3276597

and it is shown that generic lossy compression of multimodal sensory input extracts abstract representations that 1) remove modality-specific details and 2) retain information that is shared across modalities. An architecture is proposed that can learn abstract representations by identifying and retaining only the information that is shared across modalities, while discarding modality-specific information.

The benefit to real-world applications of developmental robots is often pursued through the combination of intrinsic and extrinsic motivation in a single system, an issue that is studied by the next paper [A4]. This contribution by Volpi and Polani focuses on the use of empowerment. Empowerment represents the capacity of an agent to affect its environment as an information-theoretic measure. It quantifies the ability to inject information in the environment via action and to recapture this information through sensors. The paper introduces a new formalism that combines empowerment maximization with externally specifiable, goal-directed behavior. It studies the relationship between goal-directedness and empowerment optimization, to investigate the extent to which these two behaviors can co-exist. A method to generate a behavior that is both empowered and goal-directed is proposed. This permits the design of agents capable of being "as empowered as possible" when facing an extrinsic task. The paper studies how this hybrid policy can handle problems of uncertain, or changing goals and delayed goal commitment.

The next two papers address the social aspects of motivation in the presence of humans. This issue is first examined by Rayyes et al. [A5]. A key challenge in developmental learning is the high sample complexity of algorithms. This has led to the development of intrinsic motivation approaches to drive exploration and sample gathering and thereby improve learning efficiency. However, only a few works have examined the integration of intrinsic motivation with learning from an interacting teacher. In this paper, the authors address the efficiency challenge by proposing a novel extrinsic-intrinsic motivation learning scheme. They investigate how to combine intrinsic motivation with learning from observation to accelerate learning by benefiting from human demonstration.

The second paper by Kulak and Calinon [A6] focuses on a similar combination of modalities, but uses a different approach. In this approach, a parametrized distribution of robot movement primitives is learned by combining active intrinsically motivated learning and active imitation learning. The work aims at efficiently combining experiential and observational learning. An agent can iteratively apply different learning strategies and modulate the use of these modalities based on previous experiences. The approach is demonstrated on a waste disposal task with a simulated 7-DoF Franka Emika robot. In each step of the learning process, the robot actively chooses between observational or imitation learning and experiential or intrinsically motivated learning.

Finally, three papers look at applications. Lu et al. [A7] considered how goal-oriented dialog systems use natural language accurately and efficiently to exchange information with people. The paper develops deep reinforcement learning algorithms to improve the efficiency of dialog policy learning. First, a novel "hindsight" approach is presented that uses unsuccessful dialog instances as performance feedback to the dialog learning agent. Second, user modeling is introduced to enable the agent to learn from the simulated interaction experience. Finally, a meta-learning algorithm enables the dialog agent to learn adaptively from simulated users and hindsight experience at the same time.

Jain and Kasaei [A8] examined service robots designed to work independently and adapt to the dynamic changes in the environment. In such scenarios, it is particularly important to continually learn to recognize new objects when they become available. The paper proposes a model that can learn about 3-D objects in an open-ended fashion. It uses a deep transfer learning-based approach combined with dynamically expandable layers. Experiments show that this model sets a new state-of-the-art benchmark, not only based on accuracy but also with respect to computational complexity.

Finally, Gama et al. [A9] examined how a robot can learn about its own body. Learning about one's own body is a foundational step in infant development. Spontaneous touches to the body in early infancy may give rise to early body models and bootstrap further development such as reaching skills. Unlike visually elicited reaching, reaching to one's own body requires connections of the tactile and motor space only. However, although vision is not required, the challenges of high dimensionality and redundancy of the motor system remain. This study presents an embodied computational model for learning self-touch on a simulated humanoid robot with artificial sensitive skin. The authors demonstrate that this task can be achieved on the scale of a humanoid by employing the computational frameworks of intrinsic motivations and goal babbling. Results are discussed in relation to infant studies on spontaneous body exploration as well as reaching to vibrotactile targets on the body.

III. Conclusion

In conclusion, advancing our knowledge in open-ended learning is currently at the forefront of artificial intelligence and machine learning research. Research on intrinsically motivated open-ended learning now has an unprecedented opportunity to substantially advance the state of the art. This special issue reports on progress in architectures and algorithms for intrinsically motivated learning, including architectures that combine intrinsic and extrinsic motivation; social motivation including human in the loop systems; and applications, including dialog, object recognition, and tactile exploration. It brings together the most recent state-of-the-art research in intrinsically motivated, open-ended learning for autonomous agents and developmental robots.

Appendix: Related Articles

[A1] M. I. Sener, Y. Nagai, E. Oztop, and E. Ugur, "Exploration with intrinsic motivation using object–action–outcome latent space," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 325–336, Jun. 2023.

[A2] A. Zhang, Y. Gao, Y. Niu, X. Li, and Q. Chen, "Intrinsic plasticity for online unsupervised learning based on soft-reset spiking neuron model," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 337–347, Jun. 2023.

[A3] C. Wilmot, G. Baldassarre, and J. Triesch, "Learning abstract representations through lossy compression of multimodal signals," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 348–360, Jun. 2023.

[A4] N. C. Volpi and D. Polani, "Goal-directed empowerment: Combining intrinsic motivation and task-oriented behavior," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 361–372, Jun. 2023.
[A5] R. Rayyes, H. Donat, J. Steil, and M. Spranger, "Interest-driven exploration with observational learning for developmental robots," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 373–384, Jun. 2023.
[A6] T. Kulak and S. Calinon, "Combining social and intrinsically motivated learning for multitask robot skill acquisition," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 385–394, Jun. 2023.
[A7] K. Lu, Y. Cao, X. Chen, and S. Zhang, "Efficient dialog policy learning with hindsight, user modeling, and adaptation," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 395–408, Jun. 2023.
[A8] S. Jain and H. Kasaei, "3D_DEN: Open-ended 3-D object recognition using dynamically expandable networks," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 409–418, Jun. 2023.
[A9] F. Gama, M. Shcherban, M. Rolf, and M. Hoffmann, "Goal-directed tactile exploration for body model learning through self-touch on a humanoid robot," *IEEE Trans. Cogn. Devel. Syst.*, vol. 15, no. 2, pp. 419–433, Jun. 2023.

REFERENCES

[1] S. Thrun and T. M. Mitchell, "Lifelong robot learning," *Robot. Auton. Syst.*, vol. 15, nos. 1–2, pp. 25–46, 1995.
[2] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Díaz-Rodríguez, "Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges," *Inf. Fusion*, vol. 58, pp. 52–68, Jun. 2020.
[3] D. Isele, "Lifelong reinforcement learning on mobile robots," Ph.D. dissertation, Dept. Comput. Inf. Sci., Univ. Pennsylvania, Philadelphia, PA, USA, 2018.
[4] G. Baldassarre and M. Mirolli, Eds., *Intrinsically Motivated Learning in Natural and Artificial Systems*. Berlin, Germany: Springer, 2013.
[5] P. Dhakan, K. Merrick, N. Siddique, and I. Rañó, "Open-ended continuous learning of compound goals," *IEEE Trans. Cogn. Devel. Syst.*, vol. 13, no. 2, pp. 274–285, Jun. 2021.
[6] V. G. Santucci, P.-Y. Oudeyer, A. G. Barto, and G. Baldassarre, "Editorial: Intrinsically motivated open-ended learning in autonomous robots," *Front. Neurorobot.*, vol. 13, p. 115, Jan. 2020.
[7] D. Silver, Q. Yang, and L. Li, "Lifelong machine learning systems: Beyond learning algorithms," presented at the AAAI Spring Symp., Palo Alto, CA, USA, 2013.
[8] V. Santucci, G. Baldassare, and M. Mirolli, "GRAIL: A goal-discovering robotic architecture for intrinsically-motivated learning," *IEEE Trans. Cogn. Devel. Syst.*, vol. 8, no. 3, pp. 214–231, Sep. 2016.
[9] K. Merrick and M. L. Maher, *Motivated Reinforcement Learning: Curious Characters for Multiuser Games*, R. Nugent Ed. Berlin, Germany: Springer, 2009.
[10] L. Lonini, S. Forestier, C. Teulière, Y. Zhao, B. E. Shi, and J. Triesch, "Robust active binocular vision through intrinsically motivated learning," *Front. Neurorobot.*, vol. 7, p. 20, Nov. 2013.
[11] F. Kaplan and V. V. Hafner, "The challenges of joint attention," *Interact. Stud.*, vol. 7, no. 2, pp. 135–169, 2006.

KATHRYN KASMARIK, *Guest Editor*
School of Systems and Computing
University of New South Wales Canberra
Canberra, ACT 2612, Australia
E-mail: k.kasmarik@adfa.edu.au

GIANLUCA BALDASSARRE, *Guest Editor*
Italian Institute of Cognitive Sciences and Technologies
00185 Rome, Italy
E-mail: gianluca.baldassarre@istc.cnr.it

VIERI GIULIANO SANTUCCI, *Guest Editor*
Italian Institute of Cognitive Sciences and Technologies
00185 Rome, Italy
E-mail: Vieri.santucci@istc.cnr.it

JOCHEN TRIESCH, *Guest Editor*
Frankfurt Institute for Advanced Studies
60438 Frankfurt, Germany
E-mail: triesch@fias.uni-frankfurt.de

Kathryn Kasmarik (Senior Member, IEEE) received the Bachelor of Computer Science and Technology degree (Advanced, First-Class Hons., University Medal) from the University of Sydney, Sydney, NSW, Australia, in 2002, and the Ph.D. degree in computer science from the National ICT Australia, Sydney, and the University of Sydney in 2007.

She is a Professor of Computer Science with UNSW-Canberra, Canberra, ACT, Australia. Her research lies in the areas of autonomous mental development and computational motivation, with applications in virtual characters, developmental robotics, and intelligent environments.

Gianluca Baldassarre received the bachelor's and master's degrees in economics in 1998 and a Specialization Course in "Cognitive Psychology and Neural Networks" in 1999 from the Sapienza University of Rome, Rome, Italy, and the Ph.D. degree in computer science from the University of Essex, Colchester, U.K., in 2003.

He later joined the Italian Institute of Cognitive Sciences and Technologies, National Research Council, Rome, where he has been a Postdoctoral Researcher since 2006, the Director of Research since 2021, the Coordinator of the Research Group ISTC-CNR "Laboratory of Embodied Natural and Artificial Intelligence" since 2016, the Coordinator of the EU Projects "IM-CLeVeR—Intrinsically Motivated Cumulative Learning Versatile Robots" from 2009 to 2013 and "GOAL-Robots—Goal-Based Open-Ended Autonomous Learning Robots" from 2016 to 2021, and the President of the "Advanced School in AI" since 2018. His research interests span open-ended learning of sensorimotor skills, driven by extrinsic and intrinsic motivations, in animals, humans, and robots.

Vieri Giuliano Santucci received the Bachelor of Philosophy degree from the University of Rome "Sapienza," Rome, Italy, in 2006, and the Ph.D. degree in computer science from the University of Plymouth, Plymouth, U.K., in 2016.

Since 2010, he has been a Researcher with the Institute of Cognitive Science and Technologies, Italian National Council of Research, Rome, where he got a permanent position from 2018. Since 2020, he has also been a contract Professor with Uninettuno University, Rome. His research mainly focuses on developmental robotics and machine learning, and in particular on the study of intrinsically motivated open-learning systems. He is also interested in topics related to computational neuroscience and cognitive sciences.

Jochen Triesch (Member, IEEE) received the Diploma and Ph.D. degrees in physics from the University of Bochum, Bochum, Germany, in 1994 and 1999, respectively.

After two years as a Postdoctoral Fellow with the Computer Science Department, University of Rochester, Rochester, NY, USA, he joined the faculty of the Cognitive Science Department, University of California at San Diego, La Jolla, CA, USA, as an Assistant Professor in 2001. In 2005, he became a Fellow of the Frankfurt Institute for Advanced Studies (FIAS), Frankfurt, Germany, where he is the Johanna Quandt Research Professor of Theoretical Life Sciences since 2007. He also holds professorships with the Department of Physics and the Department of Computer Science and Mathematics, Goethe University, Frankfurt. In 2019, he obtained a Visiting Professorship with the Université Clermont Auvergne, Clermont-Ferrand, France. His research interests span computational neuroscience, machine learning, and developmental robotics.

Dr. Triesch received the Marie Curie Excellence Center Award of the European Union in 2006.