Guest Editorial
Special Issue on Development and Learning and on Epigenetic Robotics

I. Scope of This Special Issue

IEEE ICDL (the Joint IEEE International Conference on Development and Learning—formerly named ICDL-EpiRob) is a unique conference gathering researchers from computer science, robotics, psychology, and developmental studies to share knowledge and research on how humans and animals develop sensing, reasoning, and actions. This includes taking advantage of interaction with social and physical environments and how cognitive and developmental capabilities can be transferred to computing systems and robotics. This approach goes hand in hand with the goals of both understanding human and animal development and applying this knowledge to improve future intelligent technology, including robots that will be in close interaction with humans.

In 2019, the conference was held in Oslo, Norway—presenting 51 papers covering a wide range of topics related to learning and development in robotic and biological systems, including language acquisition, sensorimotor development, action selection and planning, social learning, and cognitive vision. The authors of selected papers were invited to submit extended versions for this TCDS special issue. In addition, there was an open call for new submissions. The submitted papers have undergone peer review including one or more revised versions being checked by the initial reviewers. This resulting special issue presents nine of the best new submissions and submissions building on ICDL-EpiRob 2019 papers, including the winner of the best paper award.

II. Contributions to the Special Issue

Mirroring the multidisciplinary nature of the ICDL-conference, this special issue contains nine papers with a wide variety of topics related to the development and learning in humans, animals, and robots. We think they provide an important overview of state-of-the-art research. The accepted papers are grouped into three distinct categories: the first covers the studies of human development and learning; the second contains the papers on computational models mimicking human/animal development and learning; and the third contains the work on AI methods inspired by human perception and cognition.

Digital Object Identifier 10.1109/TCDS.2022.3210665

III. Category: Studies of Human Development and Learning

The first paper by Elmlinger et al. [A1] of the special issue gives new insights into early language development. It is well known that early language development in infants is affected by the infant-directed speech of their caregivers. However, the role of the infants themselves play in eliciting such caregiver speech has not been widely investigated. This paper sheds light on that topic by giving new insights into how caregivers simplify their speech in response to infants' vocalizations. Previous work has indicated that caregivers indeed simplify their speech, but it has been unclear if the effect is stable also for larger samples of talk. Through the use of a new sampling technique, the study confirms the effect and even suggests that it increases at larger word counts: speech responding to infants' vocalizations remains simplified, while noncontingent speech increases in complexity. The paper concludes by proposing two testable hypotheses to guide future work into the question of *why* caregivers simplify their speech when responding to infants' vocalizations—showing a way toward further understanding of this important part of early language development.

The second paper by Sciutti and Sandini [A2] offers new insights into the development of strategies for object exploration in children. The paper asks the question of how the exploration and recognition of an object are affected by the ability to freely move the object around—and how such exploration strategies vary across age, the complexity of the exploration task, and available senses (with or without a blindfold). An iCube, a cube with sensors measuring orientation and contact, was the target for exploration—the use of this smart device enabled efficient data gathering on exploratory strategies. The data gathered demonstrates differences between the visuo-haptic and haptic exploration strategies employed by children and adults, shedding new light on this aspect of sensorimotor development. The authors also suggest that more widespread use of tools like the iCube could in the future make assessment of exploratory strategies more efficient—reducing the need for time-consuming video annotation.

IV. Category: Computational Models Mimicking Human/Animal Development and Learning

The third paper by Philippsen and Nagai [A3] uses a computational model to suggest a new hypothesis for differences in social-cognitive processing between humans and chimpanzees.

Several studies have shown that humans and chimpanzees have very different abilities related to the social understanding of intentions. This paper hypothesizes that these differences may be explained through the theory of predictive coding: differences in how humans and chimpanzees utilize predictions when processing new information may cause the differences in behavior. As a specific task to test this hypothesis, the authors considered the problem of filling in incomplete drawings. Earlier studies have shown that human children master this at an early age, but chimpanzees instead tend to scribble randomly or follow existing lines in drawings. Recurrent neural networks were trained to fill in drawings, but with different parameters corresponding to how strongly the networks rely on their predictions during development. This computational model qualitatively replicates findings from the drawing study on humans and chimpanzees—suggesting that the ability to utilize predictions during information processing may be a key cause of differences in social-cognitive processing between humans and chimpanzees.

V. Category: Improved AI Methods Through Inspiration From Human Cognition

The fourth paper by Alamri and Pugeault [A4] takes inspiration from how humans interpret visual scenes to improve the performance of state-of-the-art object detection algorithms. When recognizing objects, we can greatly benefit from taking the objects' *context* into consideration (rather than studying the object in isolation). For instance, a colored cylinder may on its own be hard to recognize, but if we observe it on a breakfast table next to a plate, we may quickly recognize it as a cup. This paper builds on previous work on utilizing context to improve deep-learning-based object detection methods, proposing a new machine learning approach that can be integrated into most object detection methods as a post-processing step. The new method utilizes a larger number of contextual relationships than previous techniques and is demonstrated to enhance the performance of two state-of-the-art object detection algorithms. The paper shows that semantic, spatial, and scale relationships enhance object detection performance and that iterating the application of the contextual models can further enhance performance.

The fifth paper by Kerzel et al. [A5] develops and tests a shared architecture for multiple tasks related to observing and grasping objects. Learning complex visuomotor abilities for robots is challenging, as state-of-the-art machine learning methods typically require very many trials—which may be costly or infeasible in robot studies. This paper, therefore, extends a recent architecture inspired by biology and developmental psychology that integrates the learning of multiple tasks into one architecture with shared components—allowing more efficient learning. The previous work on this architecture revealed shortcomings which are here improved by adding training on auxiliary tasks: learning to identify and locate objects helps the robot also learn the primary task of grasping better.

The sixth paper by Tekülve and Schöner [A6] proposes a new model for the autonomous learning of beliefs. An intriguing difference between most modern learning systems in AI and human learning is that humans and other animals may perform single-shot learning, while machine learning systems depend on large amounts of data. An example is learning *beliefs* about the world, where animals have been demonstrated to learn and generalize from single instances of an experience. Building on the authors' previous work on modeling intentionality with dynamic field theory, this paper extends the model with the autonomous acquisition of beliefs. The model is demonstrated to be able to learn beliefs from single experiences, as well as to use beliefs and deactivate them when they no longer match the agent's experience.

The seventh paper by Jegorova et al. [A7] contains work with a robot learning a diverse set of successful behaviors for a given task, from which it can select the most suitable policy when faced with a new environment. Rather than learning a single policy, or a small fixed repertoire, their generative model for policies compactly encodes an unbounded number of policies and allows novel controller variants to be sampled. The method is demonstrated with a Baxter robot application of robust ball throwing in the presence of obstacles, as well as joint-damage robust throwing.

The eighth paper by Duran and del Pobil [A8] introduces a distance estimation approach with input from a monocular camera. That is, their method estimates a depth image by considering the optical displacement in the images induced by the different poses of the camera as a consequence of eye–head movements inspired by those involved in human fixation. It combines the images generated by these micro-movements to compute a depth map. Experiments using the Baxter robot in the Gazebo/ROS simulator are described to test and demonstrate the benefits of the approach in two different scenarios. The paper contains the evaluation of the influence of its parameters and its robustness in the presence of noise.

The ninth and final paper [A9] is concerned with direct training on a physical robot. Rayyes et al. proposed a method that drastically reduces the number of required training samples. This is achieved through applying a novel intrinsic motivation signal which combines knowledge-based and competency-based elements. Further, a new replay method is added which permits rapid online model update. The efficiency, as well as the applicability of their methods, are demonstrated with a physical 7-DoF Baxter manipulator. They show that their learning schemes are able to drastically reduce the sample complexity and learn the data-driven model online, even within a limited time frame.

The guest editors would like to thank the reviewers for their efforts in reviewing the papers and the authors of the selected papers for their positive responses to the reviewer's comments and suggestions. Also, special thanks are given to Editors-in-Chief Prof. Yaochu Jin and Prof. Huajin Tang; and the journal's Editorial Office assistants for their support of this special issue. We hope you will find the selected papers of interest and that they represent a valuable contribution to the interdisciplinary development and learning research domain.

KAI OLAV ELLEFSEN
Department of Informatics
University of Oslo
0315 Oslo, Norway
E-mail: kaiolae@ifi.uio.no

KERSTIN DAUTENHAHN
Department of Electrical and Computer Engineering
University of Waterloo
Waterloo, ON N2L 3G1, Canada
E-mail: kerstin.dautenhahn@uwaterloo.ca

JIM TORRESEN
Department of Informatics
University of Oslo
0315 Oslo, Norway
RITMO
University of Oslo
0315 Oslo, Norway
E-mail: jimtoer@ifi.uio.no

APPENDIX: RELATED ARTICLES

[A1] S. L. Elmlinger, D. Park, J. A. Schwade, and M. H. Goldstein, "Comparing word diversity versus amount of speech in parents' responses to infants' prelinguistic vocalizations," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 3, pp. 1036–1045, Sep. 2022.

[A2] A. Sciutti and G. Sandini, "To move or not to move: Development of fine-tuning of object motion in haptic exploration," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 2, pp. 366–374, Jun. 2022.

[A3] A. Philippsen and Y. Nagai, "A predictive coding account for cognition in human children and chimpanzees: A case study of drawing," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1306–1319, Dec. 2022, doi: 10.1109/TCDS.2020.3006497.

[A4] F. Alamri and N. Pugeault, "Improving object detection performance using scene contextual constraints," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1320–1330, Dec. 2022, doi: 10.1109/TCDS.2020.3008213.

[A5] M. Kerzel, F. Abawi, M. Eppe, and S. Wermter, "Enhancing a neurocognitive shared visuomotor model for object identification, localization, and grasping with learning from auxiliary tasks," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1331–1343, Dec. 2022, doi: 10.1109/TCDS.2020.3028460.

[A6] J. Tekülve and G. Schöner, "A neural dynamic network drives an intentional agent that autonomously learns beliefs in continuous time," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 1, pp. 90–101, Mar. 2022.

[A7] M. Jegorova, S. Doncieux, and T. M. Hospedales, "Behavioral repertoire via generative adversarial policy networks," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1344–1355, Dec. 2022, doi: 10.1109/TCDS.2020.3008574.

[A8] A. J. Duran and A. P. del Pobil, "Robot depth estimation inspired by fixational movements," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1356–1366, Dec. 2022, doi: 10.1109/TCDS.2020.3025057.

[A9] R. Rayyes, H. Donat, and J. Steil, "Efficient online interest-driven exploration for developmental robots," *IEEE Trans. Cogn. Devel. Syst.*, vol. 14, no. 4, pp. 1367–1377, Dec. 2022, doi: 10.1109/TCDS.2020.3001633.

ACKNOWLEDGMENT

The organizing of this special issue was supported in part by The Research Council of Norway (RCN) as a part of the projects: COINMAC "Collaboration on Intelligent Machines" under Grant 261645 and 309869; "Organising the 9th Joint IEEE International Conference of Developmental Learning and Epigenetic Robotics" in Oslo, 2019, under Grant 288150; Vulnerability in the Robot Society (VIROS) under Grant 288285; and Predictive and Intuitive Robot Companion (PIRC) under Grant 312333 and through its Centre of Excellence Scheme, RITMO, under Project 262762.

Kai Olav Ellefsen received the M.Sc. and Dr.Ing. (Ph.D.) degrees in computer science from the Norwegian University of Science and Technology, Trondheim, Norway, in 2010 and 2014, respectively.

During his Ph.D. studies, he was also a Visiting Researcher with the University of Wyoming, Laramie, WY, USA, for six months, in 2013. He worked as a Postdoctoral Researcher with the Brazilian Institute of Robotics, Salvador, Brazil, from 2014 to 2016; and a Postdoctoral Researcher with the University of Oslo, Oslo, Norway, from 2016 to 2019. Since 2019, he has been an Associate Professor with the Department of Informatics, University of Oslo. His research interests include many topics within biologically inspired artificial intelligence, including evolutionary robotics, reinforcement learning, and neuroevolution.

Dr. Ellefsen is a Board Member of the Norwegian Artificial Intelligence Society.

Kerstin Dautenhahn (Fellow, IEEE) received the Ph.D. (Dr.rer.nat.) degree from the University of Bielefeld, Bielefeld, Germany, in 1993.

She is the Canada 150 Research Chair of Intelligent Robotics with the University of Waterloo, Waterloo, ON, Canada, where she directs the Social and Intelligent Robotics Research Laboratory. Before moving to Canada in 2018, she founded and directed the Adaptive Systems Research Group, University of Hertfordshire (UH), Hertfordshire, U.K. After a couple of postdoctoral positions with GMD, Sankt Augustin, Germany, and VUB, Brussels, Belgium, she moved to the U.K. in 1997 to take up a tenured Assistant Professor (Lecturer) post with the University of Reading, Reading, U.K., before moving in 2000 to UH, where she got promoted to a Full Professor in 2003. Her main areas of research are social robotics and human–robot interaction, with a specific interested in real-world application.

Prof. Dautenhahn frequently gives keynote talks at international conferences, has several senior editorial roles for international journals, and her work is widely cited (H-index 84).

Jim Torresen (Senior Member, IEEE) received the M.Sc. and Dr.Ing. (Ph.D.) degrees in computer architecture and design from the Norwegian University of Science and Technology, Trondheim, Norway, in 1991 and 1996, respectively.

He was a Senior Hardware Designer with NERA Telecommunications, Billingstad, Norway, from 1996 to 1998 and Navia Aviation, Oslo, Norway, from 1998 to 1999. Since 1999, he has been a Professor with the Department of Informatics, University of Oslo, Oslo, where he was an Associate Professor from 1999 to 2005. He is also a Principal Investigator with the RITMO Centre for Interdisciplinary Studies in Rhythm, Time and Motion. He has been a Visiting Researcher with Kyoto University, Kyoto, Japan, from 1993 to 1994, Electrotechnical Laboratory, Tsukuba, Japan, for four months, in 1997 and 2000, and Cornell University, Ithaca, NY, USA, from 2010 to 2011. His research interests include artificial intelligence, ethical aspects of AI and robotics, machine learning, robotics, and applying this to complex real-world applications which have resulted in more than 250 scientific peer-reviewed papers in international journals, books, and conference proceedings.

Prof. Torresen is a member of the Norwegian Academy of Technological Sciences (NTVA) and the National Committee for Research Ethics in Science and Technology (NENT).