Guest Editorial Special Issue on Emerging Topics on Development and Learning

I. Scope of This Special Issue

THIS special issue will encompass state-of-the-art research on emerging topics related to development and learning in natural and artificial systems. The primary focus of this special issue is to explore the facets of development and learning from a multidisciplinary perspective by convening researchers from the fields of computer science, robotics, psychology, and developmental studies. We invited researchers to share knowledge and research on how humans and animals develop sensing, reasoning, and actions, and how to exploit robots as research tools to test models of development and learning. We expected the submitted contributions to emphasize the interaction with social and physical environments and how cognitive and developmental capabilities can be transferred to computing systems and robotics. This approach is in harmony with the dual objectives of comprehending human and animal development while leveraging this understanding to enhance future intelligent technologies, particularly for robots that will engage in close interactions with humans. Scientific production on topics, such as principles and theories of development and learning, development of skills in biological systems and robots, nature versus nurture, developmental stages, models on the contributions of interaction to learning, nonverbal and multimodal interaction, and models on active learning, has been increasing in the cognitive science literature during the last decade. In this issue, some of the submissions addressed the potential of applying these processes to robotics.

II. Contributions to the Special Issue

In the context of multiagent reinforcement learning for partially observable environments, agents rely on communication to convey information, enabling effective teamwork. Recent research has explored discrete and sparse communication inspired by human languages, but its utility in human-agent team (HAT) experiments has not been thoroughly investigated. Karten et al. [A1] examined the effectiveness of sparse-discrete methods in facilitating emergent communication, enhancing performance in both agent-only and human-agent teams. The authors introduce Enforcers, a scheme that promotes sparse communication within agent-only teams while demonstrating minimal performance loss in benchmark environments. Moreover, in experiments involving HATs, a prototype-based approach yields meaningful discrete tokens that expedite human learning of agent communication. Results from additional multiagent and human-agent teams experiments suggest that an appropriate level of sparsity reduces the cognitive load on humans, leading to superior team performance.

Drawing inspiration from the developmental and learning aspects of human infants, it is viable yet challenging for enabling artificial agents to learn language in an embodied fashion. The central issue is how to establish bidirectional associations between robot actions and language descriptions within a dynamic environment. Drawing inspiration from human language development, Özdemir et al. [A2] introduced a neural model that binds robot actions and language descriptions, using a paired variational autoencoder (PVAE) framework. The PVAE architecture is extended to include a pretrained large-scale language model, bidirectional encoder representations from transformers (BERT), enabling the model to comprehend and generate unconstrained natural language descriptions. The experiments conducted demonstrate the effectiveness of this approach, particularly in real-world scenarios with human instructions. The contributions of this work include showcasing the superiority of variational autoencoders, extending language capabilities with BERT, and shedding light on the scalability and potential applications of the proposed model.

Inspired by the language acquisition process in human infants who rely on statistical properties of phonological distributions and sensory co-occurrence cues, Taniguchi et al. [A3] tackled the challenge of unsupervised multimodal word discovery. The proposed method leverages phonological information as distributional cues and incorporates co-occurrence cues from multiple modalities, including vision, tactile, and auditory data. This is achieved through the integration of the nonparametric Bayesian double articulation analyzer (NPB-DAA) for phoneme and word discovery and multimodal latent Dirichlet allocation (MLDA) for categorizing object-related multimodal information. Experimental results demonstrate the effectiveness of the approach, especially in accurately segmenting words that describe object characteristics. Furthermore, the study explores the impact of varying the importance of co-occurrence cues on learning performance, indicating the potential of this method for more efficient word discovery.

Delving into the relatively unexplored domain of robot trust, Kirtay et al. [A4] focused on robot–robot interactions and its implications in the context of human–robot symbiotic societies. They introduce the notion that robots should trust

Digital Object Identifier 10.1109/TCDS.2023.3324805

partners capable of reducing their computational load, akin to human cognitive load, and examine this hypothesis through a series of experiments. The first set of experiments involves a robot receiving assistance from online instructors employing various guiding strategies, while the second set explores robot–robot interactions, with one robot (Pepper) aiding another (Nao) in learning a task. The study reveals that robot trust based on computational/cognitive load, within a decision-making framework, effectively guides partner selection and facilitates robot–robot scaffolding. This work underscores the significance of using computational load as an internal signal for assessing interaction partner trustworthiness.

The collaboration between humans and robots, known as human–robot collaboration (HRC), is a significant area of research in both manufacturing and household robotics. However, achieving a balance between efficiency and safety in HRC poses considerable challenges. Lyu et al. [A5] introduced an HRC pipeline that addresses these challenges by generating efficient and collision-free robot trajectories based on predictions of human arm and hand (AH) motions. The pipeline utilizes a recurrent neural network to train and predict AH trajectories using observed initial trajectory segments. This approach improves the accuracy of target estimation by combining the observed and predicted hand palm trajectories to predict the current AH motion target through Gaussian mixture models (GMMs). To ensure the safety of humans collaborating with robots, they proposed an optimization-based trajectory generation algorithm. This algorithm guarantees the avoidance of collisions and prioritizes the human's well-being during the collaborative task. To validate the effectiveness of the proposed system, experiments were conducted in a shared-workspace scenario involving human pick-and-place motions. The results demonstrate that the proposed pipeline accurately predicts the trajectory of the human AH and estimates the intended motion target early on. As a result, the task can be completed safely and efficiently.

Goasguen et al. [A6] focused on an inexperienced agent developing a subjective understanding of continuity in its sensorimotor experiences, which is crucial for sensorimotor approaches to perception. As it allows them to realistically interact with their environment, accounting for factors, such as sensory noise, motor trajectory errors, and uncertainties in the agent's internal representation of the interaction. They proposed a detailed formalization and experimental assessments of how a naive agent can utilize its uninterpreted sensorimotor flow to capture a subjective sensory continuity. This enables the agent to discover patterns or regularities in its experiences, thereby developing a sense of closeness. They also investigated the role of the agent's actions in relation to the spatial and temporal dynamics of its exploration of the environment. Building upon their previous work on sensory prediction, they extended the framework to effectively handle noisy data in the agent's sensorimotor flow.

The ability to utilize previously acquired knowledge to solve future tasks is crucial for an open-ended cumulative learning agent as it saves time and allows the agent to tackle new challenges more effectively. A lifelong learning agent should consider these policies as building blocks to solve upcoming tasks. Therefore, the agent needs to be able to combine the policies of primitive tasks to solve compound tasks, which are then added to its knowledge base. Taking inspiration from modular neural networks, Dhakan et al. [A7] proposed an approach to compose policies for compound tasks that involve concurrent combinations of disjoint tasks. Additionally, they hypothesized that learning in a specialized environment leads to more efficient learning. To test this hypothesis, they created scaffolded environments for their mobile robot-based experiments, allowing the robot to learn primitive skills. They demonstrated how the agent can combine these primitive skills to learn solutions for compound tasks. This approach reduces the overall training time required for acquiring multiple skills and results in a versatile agent capable of mixing and matching different skills. By leveraging prior knowledge and effectively combining policies, their approach enhances the learning capabilities of the agent. The ability to reuse and combine learned policies is a valuable asset in the field of robotics.

Mercier [A8] conducted their research within an educational context, with the aim of improving the teaching and assessment of 21st-century skills, such as computational thinking and creative problem solving. To achieve this, they proposed formalizing a specific activity called CreaCube, in which participants are tasked with assembling a set of robotic cubes into an autonomous vehicle, presenting an open-ended problem for them to solve. Their formalization of CreaCube is grounded in established learning science frameworks and incorporates insights from neuro-cognitive models to describe the behaviors observed in learners engaged in this activity. The chosen formalism is symbolic and aligned with upper ontologies to ensure precise and well-defined vocabularies. This facilitates effective communication between the various research fields involved, including learning science, cognitive neuroscience, and computational modeling. In addition to providing a clear specification, they proposed using available reasoners to make inferences that can guide the analysis of the data collected during the experiments. This operationalization of a creative problem-solving activity serves as part of an exploratory research action, aiming to delve deeper into understanding the cognitive processes involved. Furthermore, this study includes an effective proof of concept, demonstrating the feasibility and potential of the CreaCube activity in fostering the development of 21st-century skills.

It is a pressing issue in the realm of deep reinforcement learning to effectively explore sparse-rewards environments. Aubret et al. [A9] proposed a novel model, DisTop, which aims to simultaneously learn a diverse set of skills and enhance the performance of rewarding skills. DisTop constructs a discrete topology of the environment through an unsupervised contrastive loss, a growing network, and a goal-conditioned policy. This topology allows for the selection and learning of skills using a state-independent hierarchical policy. Importantly, DisTop is designed to be agnostic to the nature of the ground state representation, making it adaptable to various types of data, including high-dimensional binary data, images, or proprioceptive inputs. Experimental results indicate that DisTop is competitive with state-of-the-art algorithms on both single-task dense rewards and diverse skill discovery without rewards, particularly excelling in environments with sparse rewards. The article's findings highlight the potential of bottom-up skill discovery combined with dynamic-aware representation learning, offering a promising approach for addressing complex state spaces and reward settings.

In the domain of cognitive modeling, the development of holistic computational models of infant cognition represents a significant aspiration. In this context, Sagar et al. [A10] presented an innovative model known as BabyX, aiming to emulate the cognitive processes of a complete infant. This comprehensive model combines a simulated infant's brain, comprising multiple interactive modules, with a simulated body situated in a 3-D environment rendered through computer animation techniques. Notably, BabyX's simulated brain allows for real-time interaction with human users, imitating caregiver interactions. The study highlights aspects of the cognitive model related to event processing, with a focus on motion events, including how infants perceive and generate them through motor actions. Moreover, the authors introduce a framework for gathering data from real infant interactions with their parents, enabling controlled comparisons between the behavior of the simulated baby and real infants. While the evaluation presented in the article is preliminary, the ambition of this research is to create a platform for modeling various infant abilities, especially in the context of learning to cooperate with caregivers, which holds promise for understanding early cognitive development.

Human–robot collaborative assembly systems are increasingly used to enhance workplace efficiency, but they may also impose higher cognitive demands on workers. Focusing on the assessment of cognitive workload, Lagomarsino et al. [A11] proposed an online and quantitative framework to evaluate the cognitive workload induced by interactions with human operators or collaborative industrial robots employing different control strategies. This framework utilizes a low-cost stereo camera and advanced artificial intelligence algorithms for head pose estimation and skeleton tracking to monitor operators' attention distribution and upper body kinematics. Several experimental scenarios are designed to validate the online method's performance against conventional offline measurements. The results indicate the potential of integrating this vision-based cognitive load assessment into next-generation collaborative robotic technologies. Such integration could facilitate real-time monitoring of human cognitive states and enhance robot control strategies to improve human comfort, ergonomics, and trust in automation, making it crucial for the evolution of hybrid manufacturing environments. This article presents an innovative online method for assessing the cognitive workload of human operators collaborating with robots in assembly tasks. The framework employs vision-based techniques to monitor mental effort and psychological stress in real time, providing a promising approach to improve HRC in industrial environments.

In the context of human–robot interaction and the pursuit of versatile robots adaptable to various tasks, Hindemith et al. [A12] conducted a study on enhancing the teaching process of robots. The challenge at hand is to make the internal learning mechanisms of robots intuitive for all users, including nonexperts. The authors explore different user feedback types for teaching a robot a new movement skill, comparing feedback methods as star ratings on an absolute scale for single roll-outs with preference-based feedback through pairwise comparisons, coupled with optimization algorithms. Their focus is on the task of skill cup-and-ball. Results from an experimental investigation demonstrate no significant differences in subjective user experiences, but a noteworthy disparity in learning performance. The preference-based system exhibits quicker learning, although this did not affect users' evaluations of it. A follow-up study confirms the impact of human users on learning performance. This research sheds light on the importance of feedback types in teaching robots new skills and how preference-based feedback can enhance the learning process, which aligns with the broader field of development and learning.

In a world where robots are becoming increasingly integrated into human environments, the challenges in ensuring seamless human–robot interaction are of paramount importance. Duarte et al. [A13] delved into the intricacies of human-to-human handovers involving cups with varying liquid levels and textures. Their study seeks to understand the extent to which manipulation strategies are influenced by individual preferences, the presence of liquid in the cups, and the physical properties of the cups. By analyzing the kinematic data of human handovers, the authors discern two distinct levels of carefulness in manipulation. Moreover, their research reveals that the physical attributes of cups, such as fragility and deformability, significantly impact the degree of carefulness. This research takes a significant leap by applying its findings to the realm of human–robot interaction. The authors have developed a real-time robot controller, which can detect variations in the careness of human manipulation and adapt the robot's approach accordingly. This adaptive approach has profound implications for enhancing the effectiveness of robots in understanding object properties and human intentions. In essence, it empowers robots to interact with humans more intuitively and safely, which is pivotal for their broader integration into human environments.

In the realm of embodied visual navigation, a pressing issue emerges concerning the misalignment of action spaces between training and testing phases, known as embodied mismatch. Liu et al. [A14] tackled this challenge by introducing a novel visual navigation task subject to embodied mismatch. This scenario reflects real-world conditions where action spaces may differ due to various factors, such as deviations in navigation processes, robot characteristics, or unforeseen failures. To address this, the authors establish a two-stage robust adversary learning framework. In the first stage, they employ adversary training to build a robust feature representation of the agent's state. The second stage focuses on adaptation training, transferring the learned navigation strategy to a new action space with limited training samples. Experimental validation across various embodied visual navigation tasks within 3-D indoor scenes demonstrates the efficacy of their approach. The proposed framework provides a solution for learning models that can flexibly adapt to different action spaces, essential for robust embodied navigation.

The ability of robots to respond appropriately to stimuli plays a crucial role in creating the illusion of "life" in robots, thereby enhancing their acceptance as companions. Giménez et al. [A15] explored the impact of bio-inspired reactive responses on the perception of social robots. The study focused on how endowing robots with the ability to respond to stimuli, not only related to the current task but also to other external events, contributed to the illusion of life in robots. An experiment was conducted where participants observed a video-recorded interaction with two robots: one capable of responding to both task-related and nontask-related events, and the other only responsive to task-related events. The RoSAS questionnaire was used to evaluate the experiment. The results indicated significant differences in the two factors, demonstrating that the inclusion of responses to nontask-related stimuli increased the robot's perceived warmth and competence.

The significance of kinematic features in human action similarity judgments was explored by Nair et al. [A16]. The study compared the results of three experiments conducted with human participants to those of a computational model that tackles the same task. The selected model, which is rooted in developmental robotics, utilizes learned kinematic primitives to classify actions. The comparative analysis revealed that both the model and human participants are capable of accurately determining whether two actions are similar or not. Notably, most actions can be assessed for similarity using limited information from a single feature domain, such as velocity or spatial data. However, both velocity and spatial features are necessary to achieve a level of performance equivalent to that of humans in evaluating actions. The experimental findings also indicated that human participants primarily rely on kinematic information rather than action semantics when performing action identification tasks. Overall, the results demonstrated the high accuracy of both the model and human performance in action similarity judgments based on kinematic-level features, which serves as a crucial foundation for the classification of human actions.

Difficulties in establishing shared understanding during everyday conversations often arise, leading to self-repairs where speakers modify their utterances mid-turn. Previous research has primarily focused on verbal cues such as speech disfluencies, but recent studies indicate that self-repairs are also associated with specific hand movement patterns. Özkan et al. [A17] expanded on previous findings by examining the head and hand movements of both speakers and listeners, using motion parameters, such as height and 3-D velocity. The results demonstrated that speech sequences with self-repairs can be differentiated from fluent ones, as speakers exhibit higher hand and head movements, as well as increased movement speed during self-repairs. These findings, based on the analysis of 13 unscripted dialogs, have implications for the development of advanced cognitive artificial systems for natural human–machine and human–robot interactions.

Ayub and Wagner [A18] focused on the problem of few-shot class incremental learning (FSIL) in robotics, where robots need to adapt and learn with limited data in their environments without forgetting previously learned information. The paper proposed a novel framework inspired by theories of concept learning in the hippocampus and the neocortex. The framework represents object classes as sets of clusters and stores them in memory. To prevent forgetting, the framework replays data generated by the clusters of old classes when learning new classes. The proposed approach was evaluated on two object classification data sets, achieving state-of-the-art performance for class-incremental learning and FSIL. The framework was also tested on a robot, demonstrating its ability to continually learn to classify a large set of household objects with limited human assistance.

Open-ended learning is a key area of research in developmental robotics and AI, with the goal of creating machines and robots that can acquire knowledge and skills autonomously, similar to how infants learn. Cartoni et al. [A19] focused on the challenges presented by the previously proposed benchmark called the "REAL competition," which aims to foster the development of open-ended learning robots. The benchmark involves a simulated camera-arm robot and requires the solution of multiple challenges, including exploration, sparse-rewards, object learning, generalization, task/goal self-generation, and autonomous skill learning. The paper also introduced the "REAL-X" architecture, which consists of different systems that progressively solve different versions of the benchmark by releasing initial simplifications. These systems utilize a planning approach that increases abstraction dynamically and employ intrinsic motivations to encourage exploration. Some of the systems achieve a high level of performance even under demanding conditions. In conclusion, the REAL benchmark is a valuable tool for studying open-ended learning in its most challenging form.

III. Conclusion

The articles presented in this special issue address important challenges for the study of emerging topics on development and learning. The sparse-discrete communication methods in multiagent reinforcement learning introduce "Enforcers" to enhance team performance. A neural model binding robot actions and language descriptions incorporate BERT for real-world scenarios. Unsupervised multimodal word discovery leverages phonological and co-occurrence cues. Additionally, an HRC pipeline is introduced to ensure efficient and safe robot trajectories. The modular policy composition is used for versatile task-solving. These studies collectively contribute to the emerging topics on development and learning, from communication and language learning to trust, safety, and skill development in both agent-only and human-agent contexts.

Appendix: Related Articles

[A1] S. Karten, M. Tucker, H. Li, S. Kailas, M. Lewis, and K. Sycara, "Interpretable learned emergent communication for human–agent teams," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1801–1811, Dec. 2023.

[A2] O. Özdemir, M. Kerzel, C. Weber, J. H. Lee, and S. Wermter, "Language-model-based paired variational autoencoders for robotic language learning," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1812–1824, Dec. 2023.

[A3] A. Taniguchi, H. Murakami, R. Ozaki, and T. Taniguchi, "Unsupervised multimodal word discovery based on double articulation analysis with co-occurrence cues," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1825–1840, Dec. 2023.

[A4] M. Kirtay, V. V. Hafner, M. Asada, and E. Oztop, "Trust in robot–robot scaffolding," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1841–1852, Dec. 2023.

[A5] J. Lyu, P. Ruppel, N. Hendrich, S. Li, M. Görner, and J. Zhang, "Efficient and collision-free human–robot collaboration based on intention and trajectory prediction," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1853–1863, Dec. 2023.

[A6] L. Goasguen, J.-M. Godon, and S. Argentieri, "From state transitions to sensory regularity: Structuring uninterpreted sensory signals from naive sensorimotor experiences," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1864–1878, Dec. 2023.

[A7] P. Dhakan, K. Kasmarik, P. Vance, and I. R. N. Siddique, "Concurrent skill composition using ensemble of primitive skills," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1879–1890, Dec. 2023.

[A8] C. Mercier, "An ontology to formalize a creative problem solving activity," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1891–1904, Dec. 2023.

[A9] A. Aubret, L. Matignon, and S. Hassas, "DisTop: Discovering a topological representation to learn diverse and rewarding skills," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1905–1915, Dec. 2023.

[A10] M. Sagar et al., "A platform for holistic embodied models of infant cognition, and its use in a model of event processing," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1916–1927, Dec. 2023.

[A11] M. Lagomarsino, M. Lorenzini, P. Balatti, E. D. Momi, and A. Ajoudani, "Pick the right co-worker: Online assessment of cognitive ergonomics in human–robot collaborative assembly," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1928–1937, Dec. 2023.

[A12] L. Hindemith, O. Bruns, A. M. Noller, N. Hemion, S. Schneider, and A.-L. Vollmer, "Interactive robot task learning: Human teaching proficiency with different feedback approaches," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1938–1947, Dec. 2023.

[A13] N. F. Duarte, A. Billard, and J. Santos-Victor, "The role of object physical properties in human handover actions: Applications in robotics," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1948–1958, Dec. 2023.

[A14] X. Liu, D. Guo, H. Liu, X. Zhang, and F. Sun, "Visual navigation subject to embodied mismatch," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1959–1970, Dec. 2023.

[A15] A. B. Giménez, E. Fernández-Rodicio, A. Castro-González, and M. A. Salichs, "Do you want to make your robot warmer? Make it more reactive!" *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1971–1980, Dec. 2023.

[A16] V. Nair et al., "Kinematic primitives in action similarity judgments: A human-centered computational model," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1981–1992, Dec. 2023.

[A17] E. E. Özkan, P. G. Healey, T. Gurion, J. Hough, and L. Jamone, "Speakers raise their hands and head during self-repairs in dyadic conversations," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 1993–2003, Dec. 2023.

[A18] A. Ayub and A. R. Wagner, "CBCL-PR: A cognitively inspired model for class-incremental learning in robotics," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 2004–2013, Dec. 2023.

[A19] E. Cartoni, D. Montella, J. Triesch, and G. Baldassarre, "REAL-X—Robot open-ended autonomous learning architecture: Building truly end-to-end sensorimotor autonomous learning systems," *IEEE Trans. Cogn. Develop. Syst.*, vol. 15, no. 4, pp. 2014–2030, Dec. 2023.

DINGSHENG LUO, *Leading Guest Editor*
School of Intelligence Science and Technology
Peking University
Beijing 100871, China
E-mail: dsluo@pku.edu.cn

ANGELO CANGELOSI
Department of Computer Science
University of Manchester
M13 9PL Manchester, U.K.
E-mail: angelo.cangelosi@manchester.ac.uk

ALESSANDRA SCIUTTI
Cognitive Architecture for Collaborative Technologies Unit
Istituto Italiano di Tecnologia
16163 Genoa, Italy
E-mail: alessandra.sciutti@iit.it

WEIWEI WAN
Graduate School of Engineering Science
Osaka University
Toyonaka 560-0043, Japan
E-mail: wan@sys.es.osaka-u.ac.jp

ANA TANEVSKA
Uppsala Social Robotics Lab
Department of Information Technology
Uppsala University
752 36 Uppsala, Sweden
E-mail: ana.tanevska@it.uu.se

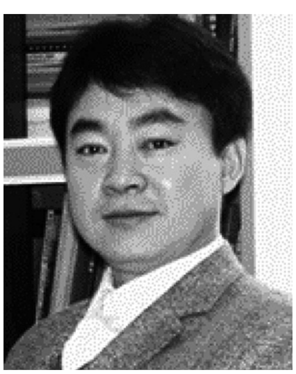

Dingsheng Luo (Member, IEEE) received the Ph.D. degree in artificial intelligence and signal processing from Peking University, Beijing, China, in 2003.

He then joined Peking University, where he has been an Associate Professor since 2005. He is currently the Vice Dean of the School of Intelligence Science and Technology and the Director of the Joint Laboratory on Intelligent Robotics and Autonomous Vehicles, Peking University. His research interests include embodied intelligence, environment perception, developmental cognitive robotics, humanoid robotics, and autonomous vehicles. He has secured over 20M RMB of research grants as a coordinator/PI and published more than 100 academic publications. He has developed with his team more than ten types of robot platforms, including PKU-HRX humanoid robot with nine models of six generations.

Dr. Luo has won more than ten academic research awards. He is a member of several IEEE Technical Committees, like Cognitive and Developmental Systems, Humanoid Robots, and Motion Control. He served as an Associate Editor for several international and domestic journals, such as IEEE TRANSACTIONS ON COGNITIVE AND DEVELOPMENTAL SYSTEMS. In 2021, he served as the General Chair for IEEE ICDL 2021.

Angelo Cangelosi received the Ph.D. degree in psychology and artificial intelligence from the Università di Genova, Genoa, Italy, in 1997.

He is a Professor of Machine Learning and Robotics with the University of Manchester, Manchester, U.K., where he is the Co-Director and the Founder of the Manchester Centre for Robotics and AI. He holds a European Research Council (ERC) Advanced Grant. He is also a Turing Fellow with The Alan Turing Institute, London, U.K. Overall, he has secured over £38 million of research grants as a Coordinator/PI, including the ERC Advanced eTALK, the UKRI TAS Trust Node and CRADLE Prosperity, the U.S. AFRL Project THRIVE++, and numerous Horizon and MSCAs grants. He has produced more than 300 scientific publications. His book *Developmental Robotics: From Babies to Robots* (MIT Press) was published in January 2015, and translated in Chinese and Japanese. His latest book *Cognitive Robotics* (MIT Press), coedited with M. Asada, was recently published in 2022. His research interests are in cognitive and developmental robotics, neural networks, language grounding, human–robot interaction and trust, and robot companions for health and social care.

Dr. Cangelosi is the Editor-in-Chief of *Interaction Studies* and *Cognitive Computation and Systems* (IET), and in 2015, he was the Editor-in-Chief of IEEE TRANSACTIONS ON AUTONOMOUS MENTAL DEVELOPMENT. He has chaired numerous international conferences, including ICANN2022 Bristol and ICDL2021 Beijing.

Alessandra Sciutti (Member, IEEE) received the Ph.D. degree in humanoid technologies from the University of Genova, Genoa, Italy, in 2010.

She is a Tenure Track Researcher and the Head of the Cognitive Architecture for Collaborative Technologies (CONTACT) Unit, Italian Institute of Technology, Genoa. After two research periods in the United States and Japan, in 2018, she was awarded the ERC Starting Grant wHiSPER (www.whisperproject.eu), focused on the investigation of joint perception between humans and robots. She published more than 80 papers in international journals and conferences and participated in the coordination of the CODEFROR European IRSES Project. The scientific aim of her research is to investigate the sensory and motor mechanisms underlying mutual understanding in human–human and human–robot interaction.

Dr. Sciutti is currently an Associate Editor of several journals, including *International Journal of Social Robotics*, IEEE TRANSACTIONS ON COGNITIVE AND DEVELOPMENTAL SYSTEMS, and *Cognitive Systems Research*.

Weiwei Wan (Senior Member, IEEE) received the Ph.D. degree in robotics from the Department of Mechano-Informatics, The University of Tokyo, Tokyo, Japan, in 2013.

From 2013 to 2015, he was a Postdoctoral Researcher with Carnegie Mellon University, Pittsburgh, PA, USA, under the support of the Japan Society for the Promotion of Science. From 2015 to 2017, he was a tenure-track Research Scientist with the National Institute of Advanced Industrial Science and Technology, Tsukuba, Japan. He is currently an Associate Professor with the Graduate School of Engineering Science, Osaka University, Toyonaka, Japan. His research interests include robotic manipulation and smart manufacturing.

Ana Tanevska received the Ph.D. degree from the University of Genova, Genoa, Italy, in 2020.

From 2020 to 2022, she was a Postdoctoral Researcher with the Center for Human Technologies, Italian Institute of Technology in Genoa, Genoa. Since October 2022, she has been a Postdoctoral Researcher with the Social Robotics Lab, Uppsala University, Uppsala, Sweden. Her research with the Social Robotics Lab is grounded in the topic of trustworthy human–robot interaction (HRI), with a particular focus on human-in-the-loop learning, and shared autonomy and agency in HRI. In addition to trustworthy HRI, her research interests include cognitive robotics (with a spotlight on the role of affect in cognition), and socially assistive HRI.