Cognitive Robotics - Towards the Development of Next-Generation Robotics and Intelligent Systems

Evi Zouganeli[1] and Athanasios Lentzas[1,2](✉)

[1] OsloMet - Oslo Metropolitan University, Oslo, Norway
{evizou,nasoslen}@oslomet.no
[2] Aristotle University of Thessaloniki, Thessaloniki, Greece

Abstract. In this paper we make the case for cognitive robotics, that we consider a prerequisite for next generation systems. We give a brief account of current cognition-enabled systems, and viable cognitive architectures, discuss system requirements that are currently not sufficiently addressed, and put forward our position and hypotheses for the development of next-generation, AI-enabled robotics and intelligent systems.

Keywords: Artificial cognition · Robotics · Intelligent systems

1 Introduction

Robots, and artificial systems more generally, are gradually evolving towards intelligent machines that can function autonomously in the vicinity of humans and interact directly with humans – e.g. drive our cars, work together with humans, or help us with everyday chores. Current artificial systems are good at performing relatively limited, repetitive, and well-defined tasks under specific conditions, however, anything beyond that requires human supervision. At the moment, it is not quite possible to deploy robots in new environments, broaden the scope of their operation, and allow them perform diverse tasks autonomously, as systems are not versatile, safe, nor reliable enough for that. Pre-programmed and pre-configured robots lack the ability to adapt, learn new tasks, and adjust to new domains, conditions, and missions.

Cognitive robotics is a multidisciplinary research field that has gained increased interest recently as it has become apparent that an advanced system architecture is a prerequisite for progressing from specialized "caged" systems to real-life autonomous systems [10]. Cognition encompasses the mental functions by which knowledge is acquired, retained, and used: perception, learning, memory, and thinking [25]. In humans, it encompasses processes such as judgment and evaluation, reasoning and computation, problem solving and decision making, comprehension, and production of language.

In order to realize such functionality in artificial systems, one needs to define an architecture that describes and governs these processes. Such system architectures are inspired by human cognition. They comprise the necessary modules

E. Zouganeli et al. (Eds.): NAIS 2022, CCIS 1650, pp. 16–25, 2022.
https://doi.org/10.1007/978-3-031-17030-0_2

for taking care of individual processes at many levels, and for overall system operation, as well as define the way information flow takes place for knowledge acquisition, reasoning, decision making, and detailed task execution. Ideally, a cognitive robot shall be able to abstract goals and tasks, combine and manipulate concepts, synthesize, make new plans, learn new behaviour, and execute complex tasks - abilities that at the moment only humans acquire, and lie in the core of human intelligence. Cognitive robots shall be able to interact safely and meaningfully and collaborate effectively with humans. Cognition-enabled robots should be able to infer and predict the human's task intentions and objectives, and provide appropriate assistance without being explicitly asked [24].

In this article we present work in progress, and our approach to cognitive robotics for next-generation systems. Our approach builds on two hypotheses/positions: i) Artificial Intelligence requires a robust cognitive architecture in order to become intelligent enough to be deployed in real-life systems in the vicinity of humans – interacting safely and meaningfully, and collaborating with humans. ii) Artificial cognitive systems need to encompass some of the processes of the right hemisphere of the human brain - such as holistic evaluation, holistic perception, intuition, imagination, and moral evaluation and reasoning.

We elaborate on these in this paper that is organized as follows. Firstly, we give an account of current cognition-enabled systems in Sect. 2. In Sect. 3 we outline a selection of cognitive architectures, and then proceed to presenting our approach and positions in Sect. 4. Finally, we conclude in Sect. 5.

2 Cognition-Enabled Robotics

Artificial cognitive systems are nowhere near human cognition at the moment, however, isolated narrow-scope cognitive functionality has been implemented in robotic systems to enable their operation. Cognition can be visualized as a pyramid [40] (Fig. 1) that models the flow of sensory input and information to realise cognitive functions and processes. The main cognitive processes are [3]: *Attention*, *Language*, *Learning*, *Memory*, *Perception*, *Thought*, and *Emotion*. Simpler processes, mostly related with behavioral elements closest to the sensory input, are at the base of the pyramid. As we move towards the top of the pyramid, more advanced and complex cognitive processes are found.

Perception is important for cognition as it provides agents with relevant information from their environment. A plethora of sensors are exploited in current systems, ranging from sensors simulating human senses (cameras, microphones etc.) [7,11], to ambient sensors and IoT devices [9]. Beyond simple object recognition, advanced perception attempts to analyze the whole scene and reason on the content of the scene [31]. Scene understanding has been used for knowledge acquisition in ambiguous situations [23].

Language-based cognitive capability has been shown to promote interaction, communication and understanding of abstract concepts [16]. Robots able to express thoughts and actions allow a better cooperation with humans [44]. An agent with the ability to summarize its actions and gain new knowledge has been demonstrated [14].

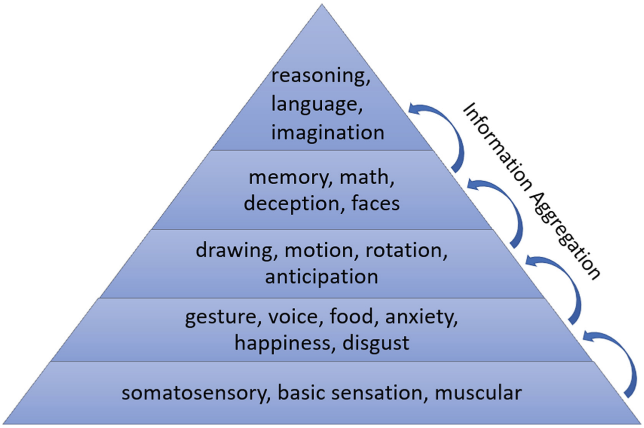

Fig. 1. Objective pyramid of cognition [40].

Learning is the core function of a cognitive system [34]. Agents can learn from expert demonstration through Imitation Learning [17], an approach that is under development. Transfer Learning is another common approach that also allows training in a simulated or protected environment [22]. Learning is currently closely woven with sensory-motor inputs and outputs, data processing, and perception, hence primarily limited to the lower layers of the cognition pyramid (Fig. 1).

The pinnacle of cognition is thinking, reasoning, decision making, planning. Reactive architectures are part of higher cognition as they affect the decision and thought process [45]. Planning and decision-making can benefit from cognition-enabled agents. Reasoning on a recognized scene allows robots to calculate an optimal path by accurately localizing itself, the goal and obstacles or dangerous areas [30]. Safety rules applied on a robot and the ability to recognize areas of potential hazard, promote a safe environment both for the robot and the humans [43]. A holistic approach to thinking with human-like cognitive reasoning and decision making processes, is far from realised, and thought processes are relatively basic at the moment.

Social robots can greatly benefit from emotional cognition [16]. Robots with the ability to recognize and express emotions (anthropomorphism) promote an easier and more effective interaction with humans [38], and robots that express empathy have been shown to help humans alter negative feelings to positive ones [5,21].

3 Cognitive Architectures

Modeling human cognition has led to the formal definition of cognitive architectures. Although first order logic approaches [20] allowed the gradual refinement of the performed actions, agents continued to lack the ability to merge new

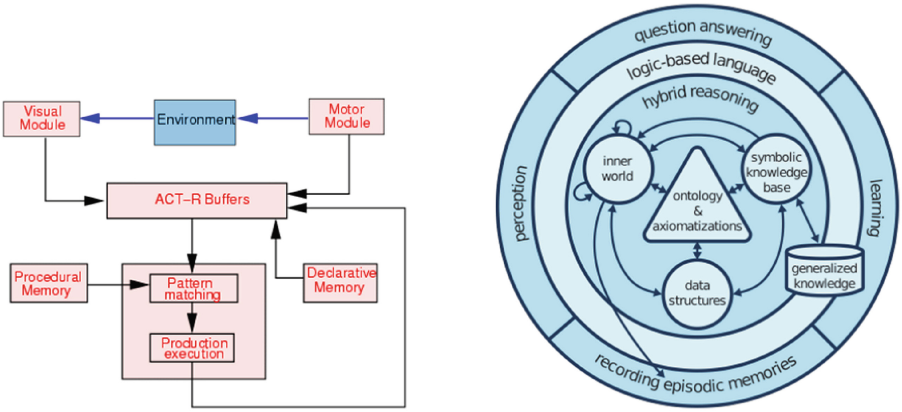

(a) ACT-R [2]   (b) KnowRob 2.0 [4]

Fig. 2. A schematic of ACT-R (a) and KnowRob 2.0 (b) architectures.

information with existing beliefs. This led to the proposal of more complex architectures. A selection of often used cognitive architectures is briefly introduced here (Fig. 2).

A commonly used architecture is ACT-R [2] where knowledge is divided based on the type of information (facts or knowledge on how to do things). Each component is accessed via a dedicated buffer, and the contents of these buffers represent the state of the world. ACT-R is based on productions, i.e. "IF" - "THEN" rules. When the current state of the world matches the precondition (using a pattern matcher module), the rule is triggered executing the relevant action. Productions, when executed, alter the state of the buffers and hence the state of the system.

A more detailed representation of human cognition is attempted by LIDA (Learning Intelligent Distribution Agent) cognitive architecture [18,19]. LIDA assumes that cognition functions on cycles with distinct phases. The first phase is perception and understanding allowing the agent to perceive the world and update the understanding of the current state. The next phase is the attention phase, where information is filtered, and the conscious content is broadcasted, followed by the action and learning phase.

The KnowRob 2.0 architecture [4] is designed specifically for robots, allowing them to perform complex tasks. At the core of the architecture are the ontologies (a subject's properties and relationships) and axioms (rules a priori true). A photorealistic representation of the environment is used for reasoning, allowing the agent to simulate its actions. Actions are stored as episodes allowing recall or knowledge transfer.

Several cognitive architectures can be considered for artificial cognition, and are extensively studied and presented by BICA [1]. In addition to the above architectures, SOAR [26], Icarus [27], and Clarion [39] are often used.

4 Position

Artificial cognitive architectures try to imitate human cognition - the epitome of cognitive systems. Some of the cognitive architectures – such as ACT-R, SOAR, LIDA – are primarily an attempt to model human cognition; whereas others – e.g. KnowRob – are inspired by human cognition but aim primarily at an architecture for artificial cognition. Cognitive architectures are progressing and gradually moving closer to human cognition, however, there is still huge uncharted ground, and a long way to go.

Semantic scene understanding, and holistic perception are only to a very basic extent realised thus far, merely at a proof-of-concept level, and there is considerable scope for further development in this area.

The importance of language in cognition was identified in early studies. Cognitive structures and capabilities are affected by language [8,37]. Despite the huge advances in speech analysis, translation, and synthesis, language is currently merely incorporated as an input/output interface in robotic systems, and is hardly included in any of the artificial cognitive processes [14,44].

Emotions have only recently been recognized as a part of cognition in humans [28,32,41] as they have previously been considered as innately hardwired into our brains. In LIDA, emotions are expressed as nodes that when triggered lead to experiencing the corresponding emotion. This is important in particular for good interaction between artificial systems and humans [13,38]. However, emotions are not incorporated in the thought process in any of the architectures or implementations, whereas in humans they often play a central role in decision making.

Currently robots are not explicitly ethical, and lack moral judgement. Ethical and moral rules have been used to that end as they can potentially affect both the acceptance of robotic applications and robotic decision making [29,33]. Norm violation may decrease human trust in an agent, therefore the agent should alter or completely discard a plan if it goes against moral values [6,12]. A fair amount of work has been done on moral reasoning and logic [15,42]. Nevertheless, moral reasoning and evaluation is not yet incorporated in cognitive architectures, neither is it an integral part of a holistic decision process. Although ethics and moral values may not be considered as part of cognition directly, in fact they play an important role in human decision making, govern human behavior, and will be instrumental for developing responsible robots.

Another relatively neglected area is artificial curiosity and imagination. While KnowRob 2.0 implements a basic form of imagination to anticipate outcomes as robots imagine the effect of their actions in their inner world representation, it is only associated to sensory-motor action and planning. Innate curiosity for exploration, global optimization, and knowledge acquisition is not explicitly accounted for in any of the reviewed architectures. This ability is critical for robots operating autonomously in unknown environments, and will allow them to effectively solve tasks even when their knowledge is not complete, and there is no human to provide the necessary information [35,36].

Moreover, current cognitive systems do not explicitly account for ingenuity. Ingenuity is the ability to employ tools or existing knowledge and use them to solve new problems in new unrelated domains. This will require complex abstraction, and synthesis of knowledge and skills. This ability will enable artificial agents to solve complex problems, and invent good solutions even when they do not have all required knowledge, sufficient experience, or the optimal tools at their disposal.

The human brain comprises two interconnected hemispheres – the left and the right – that have distinct functions and operate in different ways. The left hemisphere stands for linear thinking, detail-oriented perception, facts processing, computations, language processing, planning, logic. The right hemisphere stands for holistic thinking, holistic perception, intuitive thinking, imagination, creativity, emotional and moral evaluation. Current models of human cognition are computational in nature and represent primarily the functions of the left hemisphere. The operation and processes of the right hemisphere are by far less understood, and they are not explicitly included in the models of human cognition, let alone in robotic systems.

Our approach to attending to the above challenges in order to develop next generation robotics and intelligent systems, builds upon two main hypotheses/ positions:

i) Artificial Intelligence requires a robust cognitive architecture in order to be deployed in real-life autonomous systems in the vicinity of humans - interacting safely and meaningfully, and collaborating with humans. This hypothesis is not controversial as such, however, there is not enough awareness around this in the robotics community. Research and development in Robotics and intelligent systems has mainly targeted specific tasks and functionality - e.g. navigation, specific skill learning, etc. - rather than the overall systems architecture.
ii) In order to progress to the next level, artificial cognitive systems need to encompass some of the processes of the right hemisphere of the human brain – such as holistic evaluation, holistic perception, intuition, imagination, and moral evaluation and reasoning. This is a novel hypothesis, and needs to be proven. Our approach is to show the importance of this approach by demonstrating it in systems with superior performance.

5 Summary

In this paper we have made the case for cognitive robotics and presented our approach to next generation advanced systems. We have given an overview of human cognition, an account of cognition-enabled systems and the state of the art, and a brief outline of a selection of cognitive architectures that can lend themselves to artificial cognition. The validity of our approach remains to be demonstrated. Artificial cognitive systems are emerging, and currently at a rather early stage of development. In our opinion, they are the cornerstone towards next generation advanced robotics, the key to unlocking the potential of robots and artificial intelligence, and enabling their use in real-life applications.

References

1. Bica*ai. https://bica.ai/about/
2. Anderson, J.R., Bothell, D., Byrne, M.D., Douglass, S., Lebiere, C., Qin, Y.: An integrated theory of the mind. Psychol. Rev. **111**(4), 1036–1060 (2004)
3. Anderson, J.R.: The Architecture of Cognition. Psychology Press, New York (2009)
4. Beetz, M., Beßler, D., Haidu, A., Pomarlan, M., Bozcuoğlu, A.K., Bartels, G.: Know Rob 2.0 - a 2nd generation knowledge processing framework for cognition-enabled robotic agents. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 512–519 (2018). https://doi.org/10.1109/ICRA.2018.8460964
5. Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., Tanaka, F.: Social robots for education: a review. Sci. Robot. **3**(21), eaat5954 (2018). https://doi.org/10.1126/scirobotics.aat5954
6. Jackson, R.B., et al.: An integrated approach to context-sensitive moral cognition in robot cognitive architectures. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1911–1918 (2021). https://doi.org/10.1109/IROS51168.2021.9636434
7. Bonci, A., Cen Cheng, P.D., Indri, M., Nabissi, G., Sibona, F.: Human-robot perception in industrial environments: a survey. Sensors **21**(5) (2021). https://doi.org/10.3390/s21051571. https://www.mdpi.com/1424-8220/21/5/1571
8. Brown, R.W., Lenneberg, E.H.: A study in language and cognition. J. Abnorm. Soc. Psychol. **49**(3), 454–462 (1954). https://doi.org/10.1037/h0057814
9. Casagrande, F.D., Tørresen, J., Zouganeli, E.: Predicting sensor events, activities, and time of occurrence using binary sensor data from homes with older adults. IEEE Access **7**, 111012–111029 (2019). https://doi.org/10.1109/ACCESS.2019.2933994
10. Cassimatis, N.L., Bello, P., Langley, P.: Ability, breadth, and parsimony in computational models of higher-order cognition. Cogn. Sci. **32**(8), 1304–1322 (2008). https://doi.org/10.1080/03640210802455175. https://onlinelibrary.wiley.com/doi/abs/10.1080/03640210802455175
11. Charalampous, K., Kostavelis, I., Gasteratos, A.: Recent trends in social aware robot navigation: a survey. Robot. Auton. Syst. **93**, 85–104 (2017). https://doi.org/10.1016/j.robot.2017.03.002. https://www.sciencedirect.com/science/article/pii/S0921889016302287
12. Coeckelbergh, M.: How to use virtue ethics for thinking about the moral standing of social robots: a relational interpretation in terms of practices, habits, and performance. Int. J. Soc. Robot. **13**(1), 31–40 (2020). https://doi.org/10.1007/s12369-020-00707-z
13. Craig, M.J.A., Edwards, C.: Feeling for our robot overlords: perceptions of emotionally expressive social robots in initial interactions. Commun. Stud. **72**(2), 251–265 (2021). https://doi.org/10.1080/10510974.2021.1880457
14. DeChant, C., Bauer, D.: Toward robots that learn to summarize their actions in natural language: a set of tasks. In: Faust, A., Hsu, D., Neumann, G. (eds.) Proceedings of the 5th Conference on Robot Learning. Proceedings of Machine Learning Research, vol. 164, pp. 1807–1813. PMLR, 08–11 November 2022. https://proceedings.mlr.press/v164/dechant22a.html
15. Dennis, L., Fisher, M., Slavkovik, M., Webster, M.: Formal verification of ethical choices in autonomous systems. Robot. Auton. Syst. **77**, 1–14 (2016). https://doi.org/10.1016/j.robot.2015.11.012. https://www.sciencedirect.com/science/article/pii/S0921889015003000

16. Di Nuovo, A., Cangelosi, A.: Abstract concept learning in cognitive robots. Curr. Robot. Rep. **2**(1), 1–8 (2021). https://doi.org/10.1007/s43154-020-00038-x
17. Duan, Y., et al.: One-shot imitation learning. In: Guyon, I., et al. (eds.) Advances in Neural Information Processing Systems, vol. 30. Curran Associates, Inc. (2017). https://proceedings.neurips.cc/paper/2017/file/ba3866600c3540f67c1e9575e213be0a-Paper.pdf
18. Franklin, S., Madl, T., D'Mello, S., Snaider, J.: LIDA: a systems-level architecture for cognition, emotion, and learning. IEEE Trans. Auton. Ment. Dev. **6**(1), 19–41 (2014). https://doi.org/10.1016/j.bica.2016.04.003
19. Franklin, S., et al.: A LIDA cognitive model tutorial. Biologically Inspired Cogn. Architect. **16**, 105–130 (2016). https://doi.org/10.1016/j.bica.2016.04.003. https://www.sciencedirect.com/science/article/pii/S2212683X16300196
20. Gil, Y.: Learning by experimentation: incremental refinement of incomplete planning domains. In: Cohen, W.W., Hirsh, H. (eds.) Machine Learning Proceedings 1994, pp. 87–95. Morgan Kaufmann, San Francisco (1994). https://doi.org/10.1016/B978-1-55860-335-6.50019-2. https://www.sciencedirect.com/science/article/pii/B9781558603356500192
21. González-González, C.S., Violant-Holz, V., Gil-Iranzo, R.M.: Social robots in hospitals: a systematic review. Appl. Sci. **11**(13) (2021). https://doi.org/10.3390/app11135976. https://www.mdpi.com/2076-3417/11/13/5976
22. Hua, J., Zeng, L., Li, G., Ju, Z.: Learning for a robot: deep reinforcement learning, imitation learning, transfer learning. Sensors **21**(4) (2021). https://doi.org/10.3390/s21041278. https://www.mdpi.com/1424-8220/21/4/1278
23. Johnson-Roberson, M., et al.: Enhanced visual scene understanding through human-robot dialog. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3342–3348 (2011). https://doi.org/10.1109/IROS.2011.6094596
24. Kamarul Bahrin, M.A., Othman, M.F., Nor Azli, N.H., Talib, M.F.: Industry 4.0: a review on industrial automation and robotic. Jurnal Teknologi **78**(6–13) (2016). https://doi.org/10.11113/jt.v78.9285. https://journals.utm.my/jurnalteknologi/article/view/9285
25. Kihlstrom, J.F.: Unconscious cognition. In: Reference Module in Neuroscience and Biobehavioral Psychology. Elsevier (2018). https://doi.org/10.1016/B978-0-12-809324-5.21860-9. https://www.sciencedirect.com/science/article/pii/B9780128093245218609
26. Laird, J.E., Newell, A., Rosenbloom, P.S.: SOAR: an architecture for general intelligence. Artif. Intell. **33**(1), 1–64 (1987). https://doi.org/10.1016/0004-3702(87)90050-6
27. Langley, P., Choi, D.: A unified cognitive architecture for physical agents. In: Proceedings of the National Conference on Artificial Intelligence, Menlo Park, CA, vol. 21, p. 1469. AAAI Press, London/MIT Press, Cambridge (1999, 2006)
28. LeDoux, J.E., Brown, R.: A higher-order theory of emotional consciousness. Proc. Natl. Acad. Sci. **114**(10), E2016–E2025 (2017). https://doi.org/10.1073/pnas.1619316114. https://www.pnas.org/content/114/10/E2016
29. Lee, M., Ruijten, P., Frank, L., de Kort, Y., IJsselsteijn, W.: People may punish, but not blame robots. In: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI 2021. Association for Computing Machinery, New York (2021). https://doi.org/10.1145/3411764.3445284

30. Li, J., Chan, C.L., Le Chan, J., Li, Z., Wan, K.W., Yau, W.Y.: Cognitive navigation for indoor environment using floorplan. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 9030–9037 (2021). https://doi.org/10.1109/IROS51168.2021.9635850
31. Metaxas, D., Daniels, Z.: Image processing neural network systems and methods with scene understanding, December 2019
32. Moore, S.C., Oaksford, M.: Emotional Cognition: From Brain to Behaviour, vol. 44. John Benjamins Publishing (2002)
33. Müller, V.C.: Is it time for robot rights? Moral status in artificial entities. Ethics Inf. Technol. **23**(4), 579–587 (2021). https://doi.org/10.1007/s10676-021-09596-w
34. Peng, X.B., Andrychowicz, M., Zaremba, W., Abbeel, P.: Sim-to-real transfer of robotic control with dynamics randomization. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 3803–3810 (2018). https://doi.org/10.1109/ICRA.2018.8460528
35. Schmidhuber, J.: An on-line algorithm for dynamic reinforcement learning and planning in reactive environments. In: 1990 IJCNN, International Joint Conference on Neural Networks, vol. 2, pp. 253–258 (1990). https://doi.org/10.1109/IJCNN.1990.137723
36. Schmidhuber, J.: Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. Connection Sci. **18**(2), 173–187 (2006). https://doi.org/10.1080/09540090600768658
37. Seleskovitch, D.: Language and cognition. In: Gerver, D., Sinaiko, H.W. (eds.) Language Interpretation and Communication. HF, vol. 6, pp. 333–341. Springer, Boston (1978). https://doi.org/10.1007/978-1-4615-9077-4_29
38. Spatola, N., Wudarczyk, O.A.: Ascribing emotions to robots: explicit and implicit attribution of emotions and perceived robot anthropomorphism. Comput. Hum. Behav. **124**, 106934 (2021). https://doi.org/10.1016/j.chb.2021.106934. https://www.sciencedirect.com/science/article/pii/S0747563221002570
39. Sun, R.: 6 the CLARION cognitive architecture: toward the mind. In: The Oxford Handbook of Cognitive Science, p. 117 (2017)
40. Taylor, P., Hobbs, J.N., Burroni, J., Siegelmann, H.T.: The global landscape of cognition: hierarchical aggregation as an organizational principle of human cortical networks and functions. Sci. Rep. **5**(1), 18112 (2015). https://doi.org/10.1038/srep18112
41. Thagard, P.: Hot Thought: Mechanisms and Applications of Emotional Cognition. MIT Press, Cambridge (2008)
42. Tolmeijer, S., Kneer, M., Sarasua, C., Christen, M., Bernstein, A.: Implementations in machine ethics: a survey. ACM Comput. Surv. **53**(6) (2021). https://doi.org/10.1145/3419633
43. Vilchis-Medina, J.L., Godary-Déjean, K., Lesire, C.: Autonomous decision-making with incomplete information and safety rules based on non-monotonic reasoning. IEEE Robot. Autom. Lett. **6**(4), 8357–8362 (2021). https://doi.org/10.1109/LRA.2021.3103048
44. Wudarczyk, O.A., et al.: Robots facilitate human language production. Sci. Rep. **11**(1), 16737 (2021). https://doi.org/10.1038/s41598-021-95645-9
45. Zhang, W., Cheng, H., Hao, L., Li, X., Liu, M., Gao, X.: An obstacle avoidance algorithm for robot manipulators based on decision-making force. Robot. Comput.-Integr. Manuf. **71**, 102114 (2021). https://doi.org/10.1016/j.rcim.2020.102114. https://www.sciencedirect.com/science/article/pii/S0736584520303240