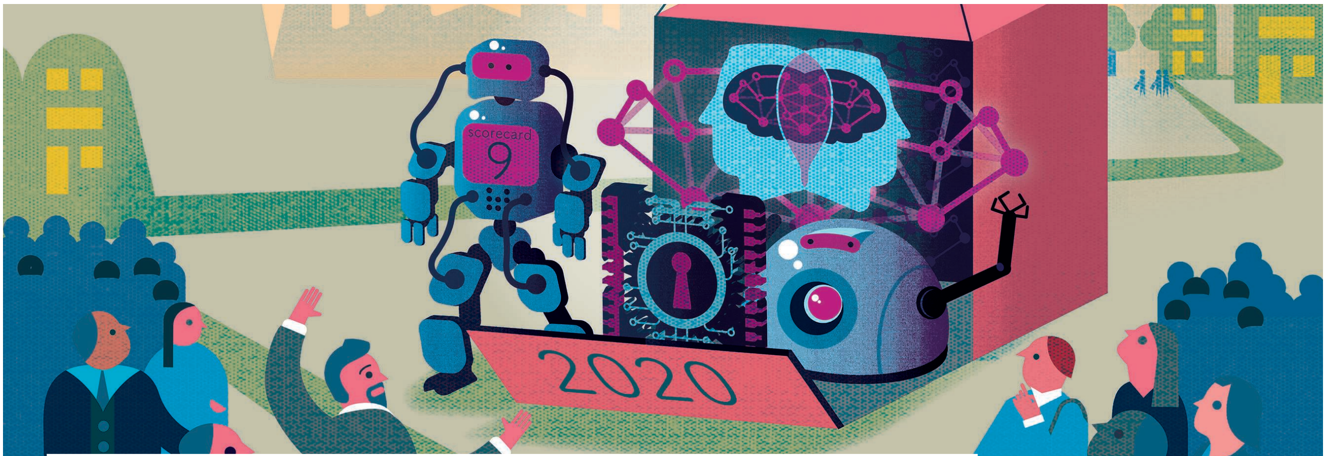

Cognitive Robotics: Making Robots Sense, Understand, and Interact

Danny Lange, Unity Technologies

Robots have spread from manufacturing floors to spaces occupied by humans. Although robots in these settings may improve the way humans work, the programming by hand of collaborative robots in such environments is increasingly difficult. We predict that recent breakthroughs in large-scale simulations, deep reinforcement learning, and computer vision collectively bring forth a basic level of cognitive abilities to robots that will lead to significant improvements of robotic applications over the next few years.

To maintain a safe environment, robotic applications traditionally restrict people from having access to the work area while the robots are active. Consequently, various uses requiring human intervention cannot be automated by robotic systems because such systems are unable to adapt to the many types of human behavior. However, improved basic cognition in robots will enable them to function in work areas previously occupied only by their human counterparts. Equipped with a sensorimotor feedback loop, including force feedback, collision detection, and computer vision, they will be safe at work alongside people.

Digital Object Identifier 10.1109/MC.2019.2942579
Date of current version: 22 November 2019

Recent advances in the field of machine learning and computer vision, including object detection, deep reinforcement learning, and imitation learning, combined with large-scale simulation in virtual reality are enabling robots to adapt more quickly to an ever-changing set of tasks. Freed from programming by hand, we predict that 2020 will be the year when progress in cognitive abilities, that is, the ability to sense, understand, and interact, will lead to new and significant improvements in robotic applications.

IMPACT

Cognition impacts numerous application domains for robotics, including pick and placement, tending, manufacturing tasks, inspection, and collaboration and assistance.

Pick and placement

Pick and placement is the task of picking a physical item from a bin and placing it in a different location. Surprisingly, despite the pervasiveness of robots in manufacturing, pick and placement remains dependent on human labor. It is also one of the most recurrent tasks performed by workers and can lead to repetitive strain injury.

Pick and placement is commonly found in the packaging process of products. Before products leave a manufacturing facility, they must be properly prepared for shipment. This may include shrink-wrapping, boxing, and packaging. Typically, these tasks are repetitive and involve small payloads, making them ideal for robots. In many industries, however, frequent product changeover makes it impossible to keep up with the reprogramming of these robots. Even innocuous product changes may require time-consuming program changes. Cognitive skills, including sensing and computer vision, allow for the quick adaptation of robots to changing production lines. Examples of additional cognitive skills include transfer learning, which enables a system to benefit from what it learned in one domain and apply it in a new domain, and continuous learning, which allows for the system to improve in perpetuity.[1]

Tending

Tending to machines is a task that is highly dependent on human labor. Operators are required to stand and watch machines for hours to address sudden operational needs. This includes tool changes or supplying raw materials. Having a robot take the tedium out of this task is challenging due to the wide range of scenarios that the robot is required to handle. Traditional robot programming does not scale to this need; however, by training a cognitive robot on a very large set of simulated scenarios, it may be able to handle and generalize common tasks to such an extent that it can safely handle unusual events never previously encountered.

Manufacturing tasks

Manufacturing tasks play an important role in production lines and are often handled by people. Examples of these tasks include gluing, dispensing, welding, and finishing. A tool is often required to carry out the task and interaction with the product can be difficult. Manufacturing tasks often require significant training for new workers.

Conventional programming of robots for handling manufacturing tasks is very challenging. The key reason is because the successful execution of these tasks inherently relies on a feedback loop involving the proper application of force, repetition, accuracy, and vision. In static settings, such as long-lasting production lines, it is feasible to automate this type of task, but the moment we face rapidly changing product dimensions and shapes, conventional reprogramming ceases to scale. A cognitive robot could be trained to master a welding technique, independent of the product's dimensions.

Inspection

A quality inspection of parts usually involves the use of high-resolution images for comparison against benchmark models. Cognitive robots are able to effectively probe objects in a rapidly changing production line without costly reprogramming. They can even learn from their own inspections and become increasingly skilled at detecting faults.

Collaboration and assistance

Without the cages and barriers that surround traditional robots, collaborative robots and people share the same workspace. One of the benefits of collaborative robots is that their precision, power, and endurance are combined with the individual skills and abilities of people.

The interactions required between humans and robots to jointly accomplish a task put immense demands on the robot's cognitive skills. Evolution has equipped humans with a wide range of tools for collaboration, including the use of language, gestures, touch, and facial expressions to facilitate interaction. Robots must support many of these communication methods to effectively and naturally collaborate with or assist a human.

Imagine a nursing robot that is tasked with feeding a patient. Not only should the robot be able to follow the patient's head movements, it would also need to understand the subtle signals that indicate when the patient is ready for the next piece of food. The robot must be able to clearly interpret the patient's voice, facial expressions, and gestures to safely carry out its task. The robot may also be required to express itself through voice, gesture, and graphical interfaces.

TECHNOLOGY CHALLENGES

We have identified a necessary paradigm shift when it comes to developing robots with cognitive skills. These systems are so intricate that they are practically impossible to program by hand. Instead, we must mimic nature and leverage the biological concept of learning from experience. For this, we use a variety of machine-learning techniques. At its core, machine learning is driven by large amounts of data, also known as *experience*. Such experience may come from real-world interactions or through simulation.

Synthetic environments

It is not always safe nor scalable to train robots in their physical environments. Aside from the obvious risks of robots destroying themselves and injuring nearby humans, wall clock time is simply too slow of a method to generate enough data within a reasonable timeframe. Physical trials are slow and costly, and the learned behaviors henceforth are very limited. Therefore, we mimic the physical environment in 3D game-like simulators, such as Unity's real-time 3D rendering platform.[2] Deployed in a highly scalable cloud infrastructure and running at a speed of several thousands of data frames per second, the equivalent of hundreds of years of experience can be generated in a matter of hours.

Simulated environments feature 3D space, time, and basic physical properties in an integrated setting, as shown in Figure 1. With the creation of static environments, like the floor, walls, and ceiling, we provide a framework for the dynamic generation of scenarios involving the robot, people, and moveable objects. By training a virtual robot in countless situations such as low-probability scenarios, it is the objective of the system to learn to generalize from these scenarios and safely handle future, yet unseen, scenarios. Furthermore, when the training is performed in a simulated environment in which core physical properties, namely, gravity, friction coefficients, and the objects' visual appearances, are randomized, it becomes apparent that the learned models successfully transfer to the physical robots despite being trained entirely in simulation.[3] This technique is also known as *domain randomization*.[4]

Procedural content generation[5] is used to systematically generate scenarios. This includes novel and random scenarios that need not be anticipated by a human knowledge engineer. The sheer scale of the simulation allows for many low-probability scenarios to be included in the data generation process. Additionally, real-world data, for instance, the sun's position, cities, weather, and traffic patterns, can be used to inform the scenario's creation.

Simulation environments can also feature a wealth of simulated sensors, including a variety of cameras and lenses, lidar, radar, and sonar. High-fidelity graphics and realistically generated signals ensure compatibility with the equivalent sensors in the physical world.

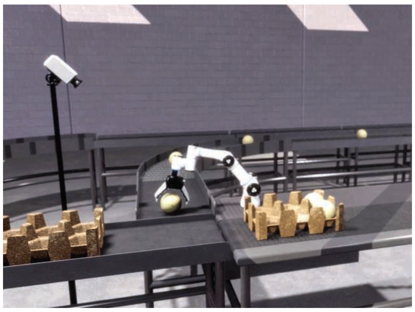

FIGURE 1. A simulated robotic environment.[6] (Source: Unity Technologies; used with permission.)

Learning methods

In cognitive robotics, it is our desire to provide robots with cognitive skills similar to those of humans and animals. As such, we must take an inclusive view of the system, including its motor and perceptual systems as well as its environmental interactions. The acquisition of knowledge through actions or perception is an important focus of cognitive robotics research.[7]

Let us use imitation learning as a starting point for teaching robots people skills. The goal of imitation learning is to mimic the human behavior used for a given task. Employing human demonstrations, a robot is trained to perform a task by learning a mapping between observations and actions.[8] Demonstrations in connections with virtual reality and motion capture permit the efficient and scalable recording of demonstration data. Having an experienced human operator provide good examples to learn from speeds up the process of learning for a performant model.

Deep reinforcement learning represents an important step toward building cognitive systems with a higher level understanding of the physical world. In the past few years, reinforcement learning has scaled to previously intractable problems, in particular, learning to play video games directly from pixels.[9] Although deep reinforcement learning is essentially learning from the simple trial-and-error method, it possesses a great capacity to learn to solve complicated tasks "from scratch" without any handcrafted rules. We refer to this as *tabula rasa*, or *learning from a clean slate*.

This learning technique is extremely data hungry, and its success has primarily been driven by the availability of large-scale simulations.

Imitation learning and reinforcement learning can be combined with curriculum learning. It is well known that people learn better when training examples are not randomly presented but, rather, systematically organized in a manner that introduces and illustrates more (complex) concepts in a gradual manner. A similar phenomenon is found in machine learning, as experiments have shown that significant improvements in learning can be achieved this way.[10] For example, to train a packaging robot, one would first train it to simply move an object from one place on a table to another. When it has learned that task, one would consecutively introduce the conveyor belt, followed by the bins, and finally, the packaging container.

There are some classes of problems, specifically, very sparse reward spaces, that become intractable for reinforcement learning with random exploration. However, recent breakthroughs in modeling human and animal curiosity has demonstrated that these limitations can be overcome.[11] It is inherently difficult to implement robots that must interact with fragile objects. A common approach has been to penalize the robot for nongentleness, which can be defined as *excessive impact force*. However, solely penalizing the robot impairs learning in a significant way because the robot will then avoid all contact with the environment. We have seen approaches that employ curiosity and deep reinforcement learning to train models that are gentle during exploration and task execution. Based on the predicting forceful contacts, the concept of curiosity has a further benefit: It encourages exploration the same way that children engage in physical risk-seeking play to probe their boundaries.[12]

Machine theory of mind

Because robots and people will work in close proximity to each other, we must draw inspiration from nature as to how people successfully interact with each other. People's ability to predict the intentions of others is essential for productive social interactions. Understanding the behavior of other people is a very important skill. Developments in neuroscience suggest that distinct regions of the brain encode personality traits that have evolved through long-term evolution and that the brain combines these traits to represent individuals with whom we meet and interact. The brain then uses this personalized model to predict the behavior of people in novel situations.[13]

Recent research has demonstrated the aptitude to train a machine to build such models too. These systems use metalearning to build models of the agents it encounters, from observations of their behavior alone.[14] The capacity for learning rich models of others will improve decision making in complex, collaborative robotic systems. This is an area ripe for new developments due to the varying definitions of cognition and the inherent intricacy of the human cognitive system, whose workings are not yet fully understood.[15]

Safety

The International Organization for Standardization (ISO) published a standard[16] that provides guidance for collaborative robot operation where a robot system and people share the same workspace. In such operations, the integrity of the safety control system is important, particularly when system parameters, for instance, the speed and force, are controlled. According to the ISO standard, a comprehensive risk assessment is required to assess not only the robot system itself but also the environment in which it is placed.

Having safety standards and regulatory oversight in place is important for the widespread adoption of cognitive robotics applications. For people to effectively interact with these devices, they must trust them. For a patient to be nursed by a robot, he or she cannot fear it. Regulatory safety standards also provide guidance to developers of cognitive robots and will provide them with a cross-industry safety framework to operate within.

Hardware

In the most recent decade, the robotics field has observed increased diversity in the hardware technology available for robotics applications beyond the traditional domain on factory floors. Versatile robotic arms with grippers and suction cups, notably, the desktop-sized UR5 arm from Universal Robots,[17] are making their way into manufacturing (see Figure 2). Among the more advanced examples of robotic hardware are human-like robotic hands, for example, the Dexterous Hand from the Shadow Robot Company.[18] Other areas of rapid development include a wide variety of sensors, such as pressure, vision/imaging, lidar, sonar, and radar. The reduced cost of acquisition, ownership, and maintenance, combined with ease of use, will accelerate the development of cognitive robotics.

Software

The open source software movement has helped spur research in robotics as well. One notable effort is the robot operating system (ROS), which helps

FIGURE 2. A desktop-sized UR5 arm.[17] (Source: Universal Robots; used with permission.)

create robot applications.[19] ROS is not actually an operating system but, rather, a middleware framework for designing robot software. It consists of a collection of tools, libraries, and conventions, with the objective of simplifying the task of creating complex and robust robot behavior across a wide variety of robotic hardware platforms.

Gazebo is an open source robot simulator that makes it possible to test algorithms, design robots, perform regression testing, and train machine-learning systems using realistic scenarios.[20] Gazebo tests robots in virtual indoor and outdoor environments and consists of a combination of a physics engine, graphical rendering, and programmatic interfaces.

RoboMaker, a cloud offering from Amazon Web Services,[21] is a solution for robot development that centers on ROS and Gazebo and removes some of the hurdles that robot developers face. There is less software to install, and the scalability of the cloud enables large-scale parallelism in simulation. Adjacent services, including video streaming, object recognition, voice command and response, and data collection, are available for integration with robotic applications.

The successful integration of middleware, simulation, and machine-learning open source packages, in combination with the scale and convenience of the cloud, may disrupt the software development process for robot applications and make cognitive robots a reality sooner rather than later.

RISKS TO PREDICTION

As discussed previously, numerous new and rapidly evolving technologies are available to aid with cognitive robotics. As a result, this prediction is not entirely without risk. We believe that risks may arise from the key areas of learning methods, model transparency, and safety and regulation.

Learning methods

Recently, we have seen considerable growth in research into machine-learning methods, where drastically new innovations are the norm, rather than the exception. Although significant resources continue to pour into the robotics field from governments, academia, and corporate initiatives, we may risk encountering presently unknown limitations of deep learning,[22] a technology that has been a significant driver behind most of the progress made in recent years.

Another area of concern is the validity of purely synthetic training data. Robotics simulation is still a young and emerging field. It may be that we cannot generate the quality and relevant training data needed to properly model the physical world. Generated scenarios may not lead to sufficient generalization for the robot to effectively interact with humans.

Model transparency

Cognitive robots use deep neural networks often described as opaque black boxes that, generally, are uninterpretable. Even if we have a complete description of all their weights, it is, in most cases, impossible for humans to even partially understand what patterns they are exploiting or know of potentially embedded flaws. Because we experience deep interaction and collaboration with cognitive robots, the need to better understand these underlying models may slow the deployment of this technology.

Robustness and predictability have been a recurring theme for robotic development; machine learning dramatically challenges that mind-set. The meaning of certification, source code auditing, and bug fixes changes as a result of this new technology. This may create adversarial situations, which can severely delay the adoption of cognitive robots.

Safety and regulation

Previously, we mentioned safety and regulatory initiatives as a positive endeavor and something that can help accelerate the development and adoption of cognitive robots. However, we may also be only one terrifying accident away from "putting the brakes" on this technology. Having robots and humans interacting and engaging in mutual physical contact carries an inherent risk. The developers of these technologies must be cognizant of safety considerations and collaborate on the methods and best practices necessary to foster deep interaction among people and robots.

Whether robots must rapidly adapt to changing production and packaging lines or engage in direct collaboration with humans, cognition is poised to thrive in the field of robotics. A number of technological areas expected to accelerate the cognitive abilities of robots was presented in this article. A shift from conventional programming to the use of large-scale simulations, deep reinforcement learning, and computer vision is likely to provide a basic level of cognitive abilities to robots, which will lead to significant improvements in robotic applications over the next few years. Regulatory safety initiatives may temporarily slow the progress in this area as may a slowdown in the previously fast-paced field of machine learning and artificial intelligence. C

ABOUT THE AUTHOR

DANNY LANGE is the vice president of artificial intelligence (AI) at Unity Technologies, San Francisco, California, where he leads efforts to advance the capabilities of AI through scalable simulations. His research interests include large-scale simulation, reinforcement learning, and robotics. Lange received a Ph.D. in computer science from the Technical University of Denmark, Copenhagen. Contact him at dlange@unity3d.com.

REFERENCES

1. S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
2. A. Juliani et al., Unity: A general platform for intelligent agents. 2018. [Online]. Available: arXiv:1809.02627
3. M. Andrychowicz et al., Learning dexterous in-hand manipulation. 2018. [Online]. Available: arXiv:1808.00177
4. J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, Domain randomization for transferring deep neural networks from simulation to the real world. 2017 [Online]. Available: arXiv:1703.06907
5. J. Togelius, G. N. Yannakakis, K. O. Stanley, and C. Browne, "Search-based procedural content generation: A taxonomy and survey," *IEEE Trans. Comput. Intell. AI in Games*, vol. 3, no. 3, pp. 172–186. 2011.
6. Unity Technologies. 2019. [Online]. Available: https://unity.com
7. L. Jurgen, "A bottom-up integration of vision and actions to create cognitive humanoids," in *Cognitive Robotics*, H. Samani, Ed. Boca Raton, FL: CRC, 2015, pp. 191–214.
8. A. Hussein, R. G. M. Medhat Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Comput. Surv.*, vol. 50, no. 2, pp. 1–35, 2017. doi: 10.1145/3054912.
9. K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, A brief survey of deep reinforcement learning. 2017. [Online]. Available: arXiv:1708.05866
10. Y Bengio, J Louradour, R Collobert, and J Weston. "Curriculum learning," in *ICML '09 Proc. 26th Annu. Int. Conf. Machine Learning*, 2009, pp. 41–48.
11. J. Achiam and S. Sastry, Surprise-based intrinsic motivation for deep reinforcement learning. 2016. [Online]. Available: arXiv:1703.01732
12. S. H. Huang et al., Learning gentle object manipulation with curiosity-driven deep reinforcement learning. 2019. [Online]. Available: arXiv:1903.08542
13. D. Hassabis, R. Nathan Spreng, A. A. Rusu, C. A. Robbins, R. A. Mar, and D. L. Schacter, "Imagine all the people: How the brain creates and uses personality models to predict behavior," *Cerebral Cortex*, vol. 24, no. 8, 1979–1987, Aug. 2014.
14. N. C. Rabinowitz, F. Perbet, H. Francis Song, C. Zhang, S. M. Ali Eslami, and M. Botvinick, Machine theory of mind. 2018. [Online]. Available: arXiv:1802.07740
15. J. N. Foerster et al., Bayesian action decoder for deep multi-agent reinforcement learning. 2018. [Online]. Available: arXiv:1811.01458
16. ISO, *ISO/TS 15066:2016: Robots and Robotic Devices—Collaborative Robots*. New York: ANSI, 2016.
17. Universal Robots. 2019. [Online]. Available: https://www.universal-robots.com
18. Shadow Robot Company. 2019. [Online]. Available: https://www.shadowrobot.com
19. Open Source Robotics Foundation. 2019. [Online]. Available: https://www.ros.org
20. Gazebo. 2019. [Online]. Available: http://gazebosim.org
21. AWS RoboMaker. 2019. [Online]. Available: https://aws.amazon.com/robomaker
22. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.