

Robots and the return to collaborative intelligence

Ken Goldberg reflects on how four exciting sub-fields of robotics — co-robotics, human-robot interaction, deep learning and cloud robotics — accelerate a renewed trend toward robots working safely and constructively with humans.

Ken Goldberg

The original robots relied on collaboration. In the play that coined the word ‘robot’ (*R.U.R.*, 1920, by the Czech writer Karel Čapek) robot workers acted collectively to rebel against unfair working conditions. And the first real robots, developed during WWII to handle radioactive materials, moved their mechanical arms under the close supervision of human ‘tele-operators’ who used levers behind shielded walls.

Since then, almost all roboticists, including me, had assumed that robots must be self-contained and carry their own power supply, memory and computing circuitry. This assumption imposed severe design constraints, limiting the ability of robots to handle uncertainty and adapt to changing conditions.

However, over the past decade robots have started to collaborate again, accelerated by advances in networking and cloud computing. Contemporary robots are immersed in a networking ecosystem that includes massive remote data centres, distributed computing, sensors, data streams and a myriad of human inputs. Robots can download data and software on demand, and perform stochastic motion planning and learning remotely both offline and online. This new generation of robots will be able to cope better with unpredictable situations and environments, and integrate usefully and safely in our world.

This Comment reviews how four growing and increasingly overlapping subfields of robotics research are influencing this trend: co-robotics, human-robot interaction, deep learning and cloud robotics.

Co-robotics

The field of telerobotics — where robots are remotely controlled by humans — has a rich history, spanning from WWII to recent advances in drones, undersea unmanned submarines, planetary rovers and surgical assist robots¹. In 1994, recognizing that a huge fraction of the cost of industrial



The Telegarden, a cloud robot that allowed anyone online to plant and water seeds in a living garden (online from 1995–2004)²³. Credit: Robert Wedemeyer

robots was spent on cages and sensors to keep humans safely away from them, a General Motors initiative sought to design a new class of safe human-assist robots, which led to the word ‘cobot’ being coined (and patented) in 1996 by J. Edward Colgate and Michael Peshkin at Northwestern University. Hami Kazerooni was doing similar work on human exoskeletons at UC Berkeley.

In 1999 Intuitive launched the da Vinci Surgical System, a minimally invasive robot fully controlled by human surgeons who work in the operating room with greatly enhanced ergonomics². This collaborative human-robot system is now used in almost a million operations per year.

In 2004, Georges Giral and Félix Ingrand co-edited a special issue of the *IEEE Robotics and Automation Magazine* on ‘human-friendly’ robots with input from researchers on several European projects. In September 2009, a group of roboticists led by Henrik Christensen published a roadmap for American robotics research³ that emphasized the importance of ‘co-robots’: robots that work directly or alongside people. This was partly motivated by the recognition that US funding agencies, specifically the US National Science Foundation (NSF), had little enthusiasm for funding technologies that could exacerbate unemployment. This led to the creation of the US NSF National

Robotics Initiative, which began in 2011 and continues today.

Pioneers of this trend introduced the term ‘collaborative robot’ to characterize robots designed to work alongside humans, with compliant joints and sensors to detect collisions and stop before humans could be harmed. In 2010 Willow Garage announced the personal robot PR2, a two-armed robot with a mobile base. In 2011 Rodney Brooks introduced a lower-cost version, the Baxter, designed to carry out a range of repetitive industrial tasks for small and medium-sized companies. This triggered a wave of research projects that explored issues around safety in robots⁴, as well as the potential of robots to learn from humans.

Since then, major robot companies FANUC, KUKA/Midea, ABB and Omron Adept have introduced collaborative robots, as have new robot companies Universal Robots, Fetch, Franka Emika and Kinova⁵. Collaborative robotics has become a fast-growing sector of the market and there is at least one industry trade show devoted to it⁶. In the past decade there has also been dramatic progress in robotic exoskeletons for assisting human workers and people with disabilities. Today, many industry labs and university research groups have initiated major projects to advance collaborative robot hardware⁷ where humans work with robots to perform assembly, inspection and warehouse order fulfilment.

Human-robot interaction

The field of human-computer interaction (HCI) originated with ‘man-machine interfaces’ in the 1960s, but it wasn’t until 2005 that the first human-robot interaction (HRI) conference was organized to focus on research into assessment and design of such interactions, reporting studies with children, adults and senior citizens interacting with humanoid robots to workers interacting with industrial robot arms. This subfield, ranging from the design of hardware to verbal and physical interactions, includes researchers in robotics, HCI, ergonomics, artificial intelligence (AI), engineering, and social and behavioural sciences. Two main areas of interdisciplinary research are systematic studies of how humans respond to robot appearance and behaviour (related to the well-known ‘uncanny valley’ effect) and the formulation of models for how robots can explicitly represent and make their intentions legible to humans⁸. HRI research seeks to optimize the collaboration between humans and robots and also addresses the next topic: how humans can actively teach robots.

Machine learning and robot learning

As is well-known, the current wave of interest in AI was sparked in 2012 with breakthrough results in computer vision enabled by ‘deep’ (many-layered) neural networks⁹. This subfield is known as ‘deep learning’ (I wish we could use the more descriptive term ‘hyper-parametric function approximation’, but it’s not nearly as catchy). Neural network and connectionist models of machine learning have a long history dating back to the 1950s. Making full use of the rise of computing power, mobile networking, availability of data and storage capabilities, deep learning was able to outperform all previous image classification algorithms in a major benchmark competition called ImageNet. Essential to its success was the availability of millions of labelled images for training, which required substantial input from humans. Fei-Fei Li, then at Princeton and later at Stanford and Google Cloud, made use of the Amazon crowdsourcing platform Mechanical Turk to incentivize and distribute human collaboration to produce the massive labelled image dataset known as ImageNet¹⁰.

Human collaboration also forms the foundation for ‘imitation learning’, where control policies are computed based on analysing video or motion-capture recordings of human demonstrations of a task such as assembly. This is an alternative to pure ‘reinforcement learning’, where robots experiment on their own to discover control policies. Imitation learning can be far more sample efficient but requires converting human demonstrations into control policies, which, given high-dimensional images as input, outputs robot control signals. There is a resurgence of interest in two variants of imitation learning: off-policy (passively observing human demonstrations) and on-policy (actively soliciting corrective feedback from humans).

By far the most active application for imitation learning is autonomous driving, where the behaviour of human drivers is recorded over millions of miles. The inputs in this case are data from cameras and lidar sensors, and the control output is the steering angle, acceleration and braking controls applied by the human driver. The ‘state space’ of an automobile based on these data is staggeringly vast.

Consider that the number of black and white low-resolution (10 pixels × 10 pixels) images is 2¹⁰⁰. This colossal number is dwarfed by the number of colour images. This means that the number of examples a driving system can learn from is an extremely minute fraction of the potential images that might be encountered in

practice. So a self-driving car must learn to generalize from an extremely limited sample. In contrast to many corporate proclamations and popular press articles, most researchers working in robotics are highly sceptical of widely cited claims that fully driverless (level 5 autonomy) vehicles will be practical in the next decade (Rodney Brooks provides many insights on his blog¹¹).

Rumour has it that certain driverless taxi systems to be released in the US in the near future will rely on a small army of human tele-operators who continuously monitor video footage from each car, ready to take control at short notice. This will require the human tele-operators to quickly assess and respond to such incidents without fatigue and on reliable, low-latency networking.

Several projects that apply deep learning to robot tasks such as grasping and assembly using ‘end-to-end’ learning, from pixels to policies, have shown promise, and the first Conference on Robot Learning (CoRL) was held in 2017. The competition to collect vast training datasets led automakers such as Tesla to begin uploading data from the drivers of all its vehicles on a daily basis. These human data are centrally used to update deep learning parameters that are periodically downloaded back to Tesla vehicles in an ongoing collaborative form of robot learning.

Cloud and fog robotics

In 1994, my students and I connected an industrial robot to the Internet, allowing anyone in the world to operate it from any browser to plant and water seeds in a living garden (pictured). The Telegarden was online 24 hours a day for nine years, and operated by over 100,000 people, more than any other robot in history¹². Building on our experience with the Telegarden in the 1990s, my lab initiated several research projects in collaborative telerobotics, where multiple human operators shared control of a single remote robot via the Internet¹³. We explored algorithms and formal models for systems like Cinematrix, where humans collaborate to track a desired trajectory¹⁴, experimenting with a system we called the Tele-Actor, where a skilled human equipped with cameras, microphones and wireless communications moves through and interacts with a difficult-to-access remote environment such as a cave, integrated circuit lab or rainforest. First-person video and audio was transmitted to a base station and then broadcast over the Internet to hundreds of online participants who interact with each other and with the remote environment by voting on goals for the human Tele-Actor¹⁵.

In subsequent years, many roboticists experimented with ‘swarm robots’ — groups of machines that interact based on common laws analogous to ants or bees — and Kiva Systems emerged in 2004, using central computing to coordinate hundreds of mobile robots in warehouses. In 2010, as the second wave of the Internet expanded rapidly with mobile phones and cloud computing, the term ‘cloud robotics’ was coined by James Kuffner, later defined as¹⁶: “Any robot system that relies on either data or code from a network to support its operation, i.e., where not all sensing, computation, and memory is integrated into a single standalone system.” Researchers including Raff D’Andrea in Europe created RoboEarth in 2009 and in 2017 companies such as iRobot and Siemens launched major cloud robotics projects. In 2018 Anki launched the Vector, the first consumer home cloud robot, and in 2019 Google will release a cloud robotics developer platform with free resources for collective map-making and object identification.

In 2012, engineers at Cisco coined the term ‘fog computing’ as an extension to cloud computing to describe systems that distribute resources between cloud-based data centres and edge devices¹⁷ to enhance performance, reliability and security. Perhaps an analogous term — ‘fog robotics’ — could describe the next generation of distributed robot systems¹⁸.

Robotics and human intelligence

The return to collaborative robotics is also consistent with several emerging subfields of research that recognize the unique ability of humans to adapt perception and control in unstructured and non-repetitive

tasks. Despite enormous progress in robot sensing, learning and control, robots cannot fully replace the unique perception and communication skills of humans.

A possible way forward is shown by new research into combinations of robots and humans that incorporate insights in collective human intelligence¹⁹. Collective intelligence builds on results in group psychology, anthropology, ecology, political science, sociology and business management theory to study and model the performance of human teams for innovation and problem solving²⁰. In contrast to a hypothetical sci-fi ‘singularity’ where superhuman AI and robot systems surpass humans²¹, I propose a constructive and inclusive alternative: ‘multiplicity’, where humans collaborate with AI and robots to mutually complement each other²².

Considered together, robots will enhance human work and life rather than replace us in our homes, hospitals, factories, farms and freeways, suggesting a future where robots are more social than solitary. □

Ken Goldberg

Department of Industrial Engineering and Operations Research, UC Berkeley, Berkeley, CA, USA.

e-mail: goldberg@berkeley.edu

Published online: 7 January 2019
<https://doi.org/10.1038/s42256-018-0008-x>

References

1. Sheridan, T. B. *Telexrobotics, Automation, and Human Supervisory Control* (MIT Press, Cambridge, 1992).
2. Rosen, J., Hannaford, B. & Satava, R. M. *Surgical Robotics: Systems Applications and Visions* Springer (Springer, New York, 2011).
3. *A Roadmap for US Robotics Research: From Internet to Robotics* (CRA, 2016); <https://cra.org/ccc/wp-content/uploads/sites/2/2016/11/roadmap3-final-rs-1.pdf>

4. Haddadin, S. et al. *Int. J. Robotics Res.* **31**, 1578–1602 (2012).
5. Tobe, F. *The Robot Report* <https://www.therobotreport.com/42-companies-empowering-robots-and-humans-to-work-side-by-side/> (2017).
6. *Vision Online* <https://www.visiononline.org/collaborative-robots-advanced-vision-and-ai-conference> (2018).
7. *US NSF National Robotics Initiative 2.0* (NSF, 2018); <https://www.nsf.gov/pubs/2017/nsf17518/nsf17518.htm>.
8. Dragan, A. D., Lee, K. T. & Srinivasa, S. S. in *8th ACM/IEEE International Conference on Human-Robot Interaction* (IEEE, 2013).
9. Beam, A. L. *GitHub* https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html (2017).
10. Gershgorn, D. *Quartz* <https://qz.com/1034972/> (2017).
11. Brooks, R. *Rodney Brooks Blog* <https://rodneybrooks.com/blog/> (2018).
12. Goldberg, K. et al. in *Proc. Second International World Wide Web Conference* (1994); <http://goldberg.berkeley.edu/pubs/Beyond-the-Web-Excavating-Oct-1994.pdf>
13. Goldberg, K., Song, D. & Levandowski, A. *Proc. IEEE* **91**, 430–439 (2003).
14. Goldberg, K. & Chen, B. in *Proc. 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2001).
15. Goldberg, K. et al. in *Proc. 2002 IEEE International Conference on Robotics and Automation* (IEEE, 2002).
16. Kehoe, B., Patil, S., Abbeel, P. & Goldberg, K. *IEEE Trans. Autom. Sci. Eng.* **12**, 398–409 (2015).
17. Bonomi, F., Milito, R., Zhu, J. & Addepalli, S. in *Proc. First Edition of the MCC Workshop on Mobile Cloud Computing* 13–16 (ACM, 2012).
18. Goldberg, K. *UC Berkeley* <http://goldberg.berkeley.edu/fog-robotics/> (2018).
19. Malone, T. W. *Superminds* (Little, Brown and Company, Boston, 2018).
20. Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N. & Malone, T. W. *Science* **330**, 686–688 (2010).
21. Kurzweil, R. *The Singularity is Near* (Viking, New York, 2005).
22. Goldberg, K. The robot-human alliance. *Wall Street Journal* (11 June 2017)
23. Goldberg, K. *UC Berkeley* <http://goldberg.berkeley.edu/garden/Ars/> (2018).

Acknowledgements

I’m grateful to the many colleagues and students I’ve learned from over the years and to those who provided specific feedback on this essay, including A. Bicchi, D. Halperin, D. Seita and R. D’Andrea.

Competing interests

The author declares no competing interests.