The Low-Cost Evolution of AI in Domestic Floor Cleaning Robots

Alexander Kleiner

■ *This article discusses AI methods deployed on domestic floor cleaning robots in the recent past and the way in which those methods are changing today. Formerly, innovations were tightly coupled with a price point customers were willing to pay. Today, there is a substantial increase in the AI found in these systems, driven by new challenges and scalable infrastructures.*

Since the early days of floor cleaning robots, manufactures have had to deal with the low-price expectations of their customers. The first models introduced in the late ’80s were equipped with state-of-the-art sensors and artificial intelligence for world modeling, collision-free path planning, and reasoning. Although those floor cleaners came with everything defining an autonomous robot today, customers were not ready to pay the price. Then, as is still too often the case today, floor cleaning robots were considered mere appliances rather than robots and thus have had to compete with their manually steered counterparts in the marketplace (Prassler, 2016).

The cleaning robots introduced during the 1990s were based on lessons learned. They came with less computational power and cheaper sensors. The first Roomba followed a very simple but effective principle: random traversal of the environment while bouncing back from walls, similar to a cue ball hitting a rail. Although this reactive behavior could not deliver guarantees on completeness, it offered a good compromise between functionality and price that satisfied people all over the world. More than 15 million of these robots have been sold.

Almost two decades after the introduction of floor cleaning technology, the first floor cleaning robots with mapping capabilities appeared. World modeling, as a prerequisite for deliberation, enabled systematic cleaning for the first time. Most of these robots deploy simultaneous localization and mapping (SLAM) techniques such as laser- and vision-based SLAM. Besides localization, cleaning robots are solving many more problems than ever before. For example, they now build maps of the environment and reason about unknown areas that need exploration. Open tasks are maintained and ranked by their expected utility, and paths are planned by A*-like methods. All this at a price point that keeps them still affordable.

After I left academe in 2014, I joined the technical organization at iRobot. I quickly learned how challenging it is to build deliberative robotic systems exposed to millions of individual homes. In contrast, the research results presented in papers (including mine) were mostly limited to a handful of environments that served as a proof of concept. Practical problems such as changing illumination and the kidnapped robot problem can have major impact on SLAM and coverage methods. Interestingly, cleaning robots are quite often exposed to the kidnapped robot problem in private homes. Consumers tend to teleport them from one place to another after they finish a task. All of these situations, and many more, are handled reliably by systematic floor cleaners today.

One limitation, however, is their inability to learn about environments. Each time the "start clean" button is pressed, an entirely new map of the environment is created and each time a mission concludes, that map is erased. This approach has the advantage that the system can easily adapt to changes in the environment over time. It has the disadvantage that there is no way for the robot to learn about potentially dangerous locations, such as environments with wires and overhanging structures it might get stuck. The user is also not able to tell the robot which areas to clean more thoroughly, and which rooms to avoid. Features such as these require a life-long representation, that is, a persistent map, of the environment and an understanding of the meaning of these locations in the human world.

With my team at iRobot, I developed techniques for structure analysis and optimized cleaning coverage based on persistent maps. The system for structure analysis was tested on several thousand grid maps of individual user homes (Kleiner, 2017). Feedback from user in-home tests helped substantially to render the interface usable for anybody. Information about structure facilitates intuitive user interfaces. For example, after an initial learning phase, robots are capable of responding to verbal commands such as "clean my kitchen." Structure information also supports optimized cleaning behaviors. Cleaning can be performed systematically, room by room, with each room treated differently with respect to its structure. With these techniques, we achieved a reduction in cleaning time of up to 50 percent as compared to the cleaning time of structure-agnostic robots. This result can theoretically still be improved by computing the optimal solution for target traversal and space coverage. To this end, one has to solve the set cover problem and the traveling salesman problem, both of which are known to be intractable on larger sets. In practice, other factors such as localization accuracy and dynamic changes over time are having a more significant impact on the strategy.

Fortunately, robots no longer have to carry all the required computational resources onboard. Thanks to cloud computing and internet connectivity, computation can be outsourced and scaled off-board. These advances in technology resolve the strong dependency between hardware costs and the computational needs of AI features. With outsourced computation, and also the lower cost of onboard sensors driven by the mobile phone industry, more sophisticated features are in reach. In the near future, floor cleaning robots will be more than just appliances. Together with other connected devices such as thermostats, cameras, and light bulbs, they will build the Smart Home. Here, the clear advantage of these cleaning robots will be that they can cover the entire home and thus will be able to learn about local changes over time. They will be able to provide sensor measurements, such as temperature and illumination, at various locations, which could potentially help other devices to optimize their specific tasks. For example, thermostats could be informed of temperature distributions, ambient illumination might be controlled by optimizing over the joint activation of several light bulbs, and much more. These new sensor and actuator modalities, along with the resulting combinatorial explosion, will raise new challenges to methods in AI and also alter the way domestic robots are seen today.

References

Kleiner, A.; Baravalle, R.; Kolling, A.; Pilotti, P.; and Munich, M. 2017. A Solution to Room-by-Room Coverage for Autonomous Cleaning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5346–5352. Piscataway, NJ: Institute for Electrical and Electronics Engineers.

Prassler, E.; Munich, M.E.; Pirjanian, P.; Kosuge, K. 2016. Domestic Robotics. In *Springer Handbook of Robotics*, edited by B. Siciliano and O. Khatib. Berlin: Springer.

Alexander Kleiner worked as a senior principal robotics scientist with technical leads at iRobot, and now works as the president of AI and Machine Learning for FaceMap, LLC. Previously, he worked as associate professor at Linkoeping University in Sweden, where he headed the research group on collaborative robotics. His research interests include collaborative robotics, navigation planning, and machine learning.