

Climbing on Giant’s Shoulders: Newcomer’s Road into the Amazon Robotics Challenge 2017

Gustavo Alfonso Garcia Ricardez^{*†1}, Lotfi El Hafi^{*†1}, Felix von Drigalski^{*†1}, Rodrigo Elizalde Zapata^{†1}, Chika Shiogama^{†1}, Kenta Toyoshima^{†1}, Pedro Miguel Uriguen Eljuri^{†1}, Marcus Gall^{†1}, Akishige Yuguchi^{†1}, Arnaud Delmotte^{†1}, Viktor Gerhard Hoerig^{†1}, Wataru Yamazaki^{†1}, Seigo Okada^{†2}, Yusuke Kato^{†2}, Ryutaro Futakuchi^{†2}, Kazuo Inoue^{†2}, Katsuhiko Asai^{†2}, Yasunao Okazaki^{†2}, Masaki Yamamoto^{†2}, Ming Ding¹, Jun Takamatsu¹, and Tsukasa Ogasawara¹

Abstract—The Amazon Robotics Challenge has become one of the biggest robotic challenges in the field of warehouse automation and manipulation. In this paper, we present an overview of materials available for newcomers to the challenge, what we learned from the previous editions and discuss the new challenges within the Amazon Robotics Challenge 2017. We also outline how we developed our solution, the results of an investigation on suction cup size and some notable difficulties we encountered along the way. Our aim is to speed up development for those who come after and, as first-time contenders like us, have to develop a solution from zero.

I. INTRODUCTION

The Amazon Robotics Challenge (ARC), formerly known as Amazon Picking Challenge, has gathered the robotics community to combine knowledge from multiple fields into practical solutions for warehouse operation. In just three years, it has become the biggest robotic challenge in the field of manipulation and warehouse automation, and has attracted international attention from both research groups and press.

Moreover, the ARC has pushed the state of the art so that the last winning team was able to perform at a speed of 100 items per hour with a success rate of nearly 84%, approaching the performance of a human worker whose full speed is 400 items per hour with an almost 100% success rate [1].

For new teams and companies developing robotic warehouse solutions, it can be hard to foresee which requirements



Fig. 1: Example of clutter with the volume-reduced storage system (end effector’s camera view).

will come into play and which failures can occur. In our own development process, we found it helpful to draw on practical knowledge from past editions of the ARC and to consider the design approaches proposed by previous contestants. In order to aid future teams and thus help speed up the emergence of new robotic solutions, we provide a summary of the problems we have encountered, how we tackled them and our successes and failures.

In this paper, we present the materials we have found that helped us start up our development effort, the lessons we learned from them, as well as the tools we were able to make use of. Finally, we discuss potential solutions to the new challenges in the 2017 edition.

The remainder of this paper is structured as follows. First, we present the new challenges and requirements of the ARC 2017 in Section II. Next, we give an overview about the resources and materials that we found from past participants, the lessons we have drawn from this and our own experiences in robotic challenges in Section III. We then outline how we developed our solution and which notable difficulties we encountered along the way in Section IV. We also present approaches and designs that we were implementing, but had to discard in later design iterations, and explain why these were not as promising as they had seemed to us. Finally, we conclude in Section V.

^{*} Authors contributed equally to the paper.

[†] Members of Team NAIST-Panasonic.

¹G. A. Garcia Ricardez, L. El Hafi, F. von Drigalski, R. Elizalde Zapata, C. Shiogama, K. Toyoshima, P. M. Uriguen Eljuri, M. Gall, A. Yuguchi, A. Delmotte, V. G. Hoerig, W. Yamazaki, M. Ding, J. Takamatsu, and T. Ogasawara are with Graduate School of Information Science, Nara Institute of Science and Technology (NAIST); 8916-5 Takayama, Ikoma, Nara 630-0192, Japan. (garcia-g, lotfi.el.hafi.kx2, felix.von.drigalski.fp6, rodrigo.elizalde.qy5, shiogama.chika.ru0, toyoshima.kenta.th2, pedro.uriguen.pl3, marcus.gall.lw3, yuguchi.akishige.xu9, arnaud.delmotte.zr3, viktor.hoerig.uw4, yamazaki.wataru.yo6, ding, j-taka, ogasawar)@is.naist.jp

²S. Okada, Y. Kato, R. Futakuchi, K. Inoue, K. Asai, Y. Okazaki, and M. Yamamoto are with Advanced Research Division, Panasonic Corporation; 1006 Kadoma, Kadoma, Osaka 571-0050, Japan. (okada.seigo, kato.yusuke001, futakuchi.ryutaro001, inoue.kazuo, asai.k, okazaki.yasunao, yamamoto.mas)@jpn.panasonic.com

II. AMAZON ROBOTICS CHALLENGE

The ARC consists of two tasks:

- Pick: move 10 items from the storage system to 3 boxes, modeling the purchase process of Amazon.
- Stow: store 20 new items in the storage system, modeling the process of adding newly arrived items into the warehouse.

The rules of the ARC 2017 [2] have significant differences compared to previous years:

- a) Half of the items are unknown until 30 minutes before the round starts.
- b) The storage system is designed by the teams.
- c) The volume of the storage system is 70% smaller than the previous years.

(a) limits the applicability of conventional learning-based approaches in which a classifier is trained with large amounts of data (*e.g.*, up to 150 000 images/item) to recognize items, which affects the approaches of eight of the top-ten teams of the ARC 2016 [3], [4]. However, this new requirement is realistic for warehouse applications, where new items are scanned and entered into the database and must be manipulated shortly after.

With (b), Amazon opens up a new design dimension in the challenge, allowing the teams to adapt the storage solution to their robot and to propose new ideas for the storage system.

As the number of items remains the same as in previous years, the reduction of volume in (c) almost inevitably causes items to be stacked and occlusions to occur, which poses a significant challenge for object recognition, manipulation and planning. Fig. 1 shows a fully stocked shelf.

These new rules require an effective system, which is more flexible and capable of working within the size constraints.

III. LESSONS FROM THE PAST

The knowledge accumulated in the two previous editions has become the baseline to build up better solutions. Although the requirements have changed for the ARC 2017, it is still instructive to look at past proposals and identify potential failures. In the remainder of this section, we show which materials gave us insights about how past teams approached the competition, what obstacles they encountered and how this knowledge shaped our own design approach.

A. Reports

Numerous reports and media coverage have summarized the state of the art of ARC, as well as the accumulated heuristics, such as [5] by Team RBO who took 1st place in ARC 2015.

Correll *et al.* [6], [7] describe platforms, grippers, sensors, and perception and motion planning techniques used by the teams competing in the 2015 edition. They conclude that there is trade-off between customization and dependability of software developed by the teams and third parties.

This was complemented by an in-depth report from Nikkei [3], [4] about the solutions of the 2016 teams. Additional reports from the Robotics Society of Japan illuminated some more approaches and problems in [8], [9].

TABLE I: Common failures and the potential impact on the performance.

Failure	Potential impact
Collision with storage system	Round loss
Planning failure	Round loss
Items left on recognition space	Object recognition capability loss
Losing suction contact	Point loss due to dropped item
Two-item grasping	Point loss due to lost item
Object recognition errors	Point loss due to misplaced item
Grasping failures	Time loss
Slow path-planning	Time loss

A number of previous competitors, such as Team C²M [10], Team R U Pracsys [11] and Team MIT-Princeton [12], also provide implementations of their approaches, as well as the datasets generated during the competitions. These datasets have been useful as a starting point to train and test our object recognition algorithms, as many of the items are also found in the ARC 2017 practice kit.

B. Interpretation

Looking at the past competitions, it becomes clear that the usage of suction cups and deep learning has tended to increase teams' success rates. Furthermore, the reports show that teams using a single robot manipulator perform better, and make a strong case for reusability by using ROS¹.

We saw in previous ARC editions that failures are unavoidable and have to be planned for. We have identified the most common problems that have occurred during the competition and summarized their potential impact in Table I.

With these failures in mind, we drew the following main conclusions to guide our development effort:

- Suction is an effective grasping tool, as 80% of the items are suctionable.
- A professional suction system is important for reliable operation.
- Learning-based object recognition can yield up to 90% success rate [3], [4].
- Using depth information does not improve object recognition significantly, and may even be counterproductive.
- Robust error recovery is fundamental for a competitive performance.
- A 7-DOF manipulator can save time by achieving the target pose quicker as demonstrated by the previous winners [5], [13].
- Task planning using state machines is effective [14].
- Modifying the code in the last minute must be avoided, as it leads to human error.
- Sensors can overheat and stability issues should be anticipated.
- Illumination in the venue significantly affects the object recognition performance.

IV. DEVELOPMENT

Our previous experience in different robotics competitions, such as the Airbus Shopfloor Challenge 2016 [15], has led us to focus on simplicity and reliability.

¹Robotic Operating System: <http://www.ros.org/>.

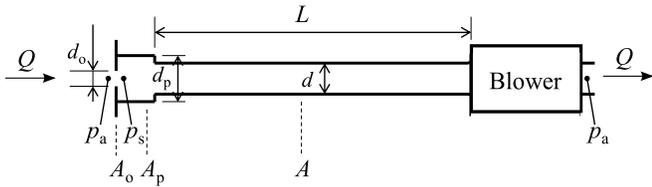


Fig. 2: Suction force model. Q is the flow rate, p_a and p_s are the external and internal static pressures, respectively. A , A_p and A_o are the cross sections of the tube, the suction cup chamber and the suction cup border, respectively, and their corresponding diameters are denoted with d . L is the length of the hose.

As newcomers whose computer vision and machine learning expertise has not been tested in an ARC competition, we decided to focus on a simple but reliable design, and to improve upon a minimal working prototype.

In this section, we present our ongoing work on the suction tool and storage system, along with the experimental results we analyzed.

A. Suction cup size and force

In the past, many teams have struggled with designing an adequate suction mechanism with sufficient suction force. Commercial vacuum cleaners and similar solutions do not generate the necessary flow and pressure difference to secure all items. On the other hand, excessive suction force may damage the packaging of the item, *e.g.*, clothes in PVC bags. To respond to this problem and to design a system that can suction all items safely, we have investigated the suction force systematically.

First, we modeled the suction tool as shown in Fig. 2 as a long tube with an opening at the end. Then, with the pressure difference and flow rate of the blower, the hose diameter, the suction cup size, and the relative opening at the end (assuming an imperfect seal), we calculate the resulting normal force.

We performed preliminary experiments with suction cups of 30 mm, 40 mm and 50 mm in diameter d_p , and hoses of 10 mm, 20 mm, 30 mm, 40 mm, and 50 mm in diameter d and 5 m in length L . Fig. 3 shows the results.

The combination of $d_p = 40$ mm and $d = 30$ mm had the best performance when tried with all items from the ARC 2017 practice kit: 36 items can be suctioned (90%). However, 9 items can be potentially damaged (22.5%), thus additional suction force control is required. Thin and cylindrical items (*e.g.*, wine glass, toilet brush, dumbbell) and the mesh cup have proven to be the most challenging. Nonetheless, it is notable that even the marbles and the body scrubber could be manipulated reliably.

The main conclusion is that with sufficient hose diameter and air flow, items can be held even if the suction seal is imperfect, such as when the item surface is uneven or rough. With smaller hose sizes, suction cups of smaller diameters break away significantly faster, as they can build

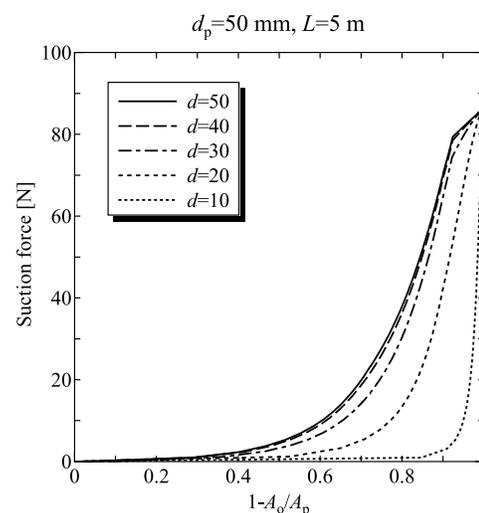
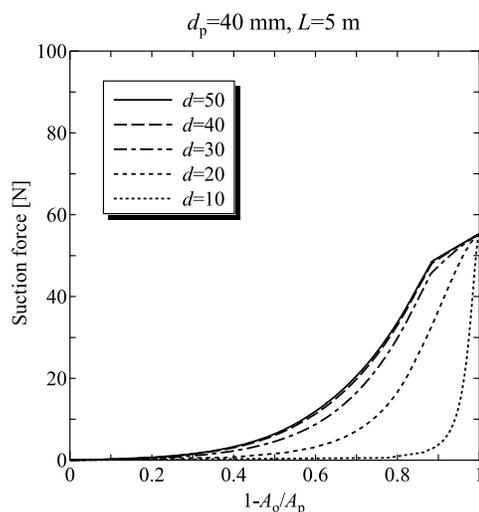
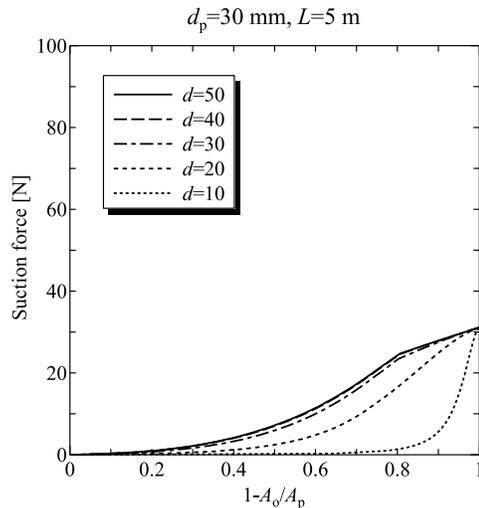


Fig. 3: Effective suction force for an imperfect suction seal, assuming the vacuum machine Induvac VC355-720 used by Team Delft in 2016. $1 - A_d/A_p$ goes to 1 for a perfect suction seal, and to 0 when no part of the object is in contact. It can be seen that the flow enabled by larger hose diameters makes the suction more robust to seal leakage and can work even for some porous objects.

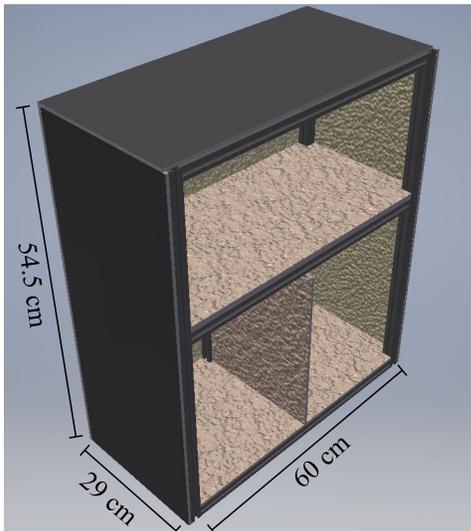


Fig. 4: Initial design of the storage system.

up a pressure differential, but cannot transport enough air to sustain a leaking seal.

It is worth mentioning that, in our first design, we used two DOFs in the suction part so that we can freely orient the item that was suctioned. We abandoned the idea after realizing that it would not yield much flexibility and advantages compared to simpler solutions, since we already use a 7-DOF arm. It is also not feasible when a wider hose is used.

B. Storage system

We started our development with a storage system composed of adjustable aluminium profiles. In retrospect, the ability to adjust division sizes was unnecessary, as we did not end up testing many different configurations. For the first prototype, we fixed the depth of the shelf to accommodate the second largest dimension of the item maximum size, *i.e.*, 27 cm. The other two dimensions were set to an approximated square, making the shelf compliant with the rule of 95 000 cm³ maximum volume. The resulting storage system is shown in Fig. 4.

During the preliminary tests with this storage system, we realized that it gets extremely cluttered, and requires very dexterous manipulation to avoid items falling out. As of now, we are experimenting with different configurations and drawers to decrease the clutter. While drawers increase the surface area of the shelf and can thus reduce clutter, they require additional manipulation, making them a potential point of failure. Moreover, the drawer volume cannot be used for large items.

We are also planning to use weight sensors underneath our storage system, totes and boxes to account for falling items and to aid during the recognition of picked items.

V. CONCLUSION

The practical knowledge from previous editions of the ARC warn us about usually neglected issues which have proved to have a critical impact in the performance. In this

paper, we presented a literature survey on past editions of the Amazon Robotics Challenge as well as the tools and resources that are available to teams and companies now entering the field. Further, we summarized which problems past contestants have faced, both in their development as well as in the competition. We extracted design guidelines and best practices, which meet the new requirements of the ARC 2017 and revolve on the concept of simplicity toward reliability. Finally, we presented our preliminary storage system and an analysis of suction force for different hose diameters and suction cup sizes, and concluded that larger hose diameters have a significantly stabilizing effect on the connection between suction cup and item.

REFERENCES

- [1] E. Ackerman, "Team Delft Wins Amazon Picking Challenge," *IEEE Spectrum*, 2016. [Online]. Available: <http://spectrum.ieee.org/automaton/robotics/industrial-robots/team-delft-wins-amazon-picking-challenge>
- [2] Amazon, "Amazon Robotics Challenge," 2016. [Online]. Available: <https://www.amazonrobotics.com/#/roboticschallenge/rules>
- [3] Nikkei, "Amazon Picking Challenge 2016 In-Depth Report (Part 1)," *Nikkei Robotics*, no. 14, pp. 4–16, 2016.
- [4] Nikkei, "Amazon Picking Challenge 2016 In-Depth Report (Part 2)," *Nikkei Robotics*, no. 15, pp. 16–22, 2016.
- [5] C. Eppner, S. Hofer, R. Jonschkowski, R. Martin-Martin, A. Sieverling, V. Wall, and O. Brock, "Lessons from the Amazon Picking Challenge: Four Aspects of Building Robotic Systems," in *Proceedings of Robotics: Science and Systems (RSS)*, 2016, pp. 36–45.
- [6] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Lessons from the Amazon Picking Challenge," *arXiv Robotics*, no. 1601.05484, pp. 1–14, 2016.
- [7] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Analysis and Observations From the First Amazon Picking Challenge," *IEEE Transactions on Automation Science and Engineering*, no. 99, pp. 1–17, 2016.
- [8] K. Okada, "Perspective from International Robotics Competitions," *Journal of the Robotics Society of Japan*, vol. 35, no. 1, pp. 9–12, 2017.
- [9] Y. Domae, "Technology Trends and Future about Picking and Manipulation by Robots," *Journal of the Robotics Society of Japan*, vol. 35, no. 1, pp. 13–16, 2017.
- [10] MPRG, "Team C²M: APC RGB-D+PointCloud Dataset 2015," 2015. [Online]. Available: http://mprg.jp/research/apc_dataset_2015_e
- [11] C. Rennie, R. Shome, K. E. Bekris, and A. F. De Souza, "A Dataset for Improved RGBD-based Object Detection and Pose Estimation for Warehouse Pick-and-Place," *arXiv Computer Vision and Pattern Recognition*, no. 1509.01277, pp. 1–7, 2015. [Online]. Available: http://pracsyslab.org/rutgers_apc_rgbd_dataset
- [12] A. Zeng, K.-T. Yu, S. Song, D. Suo, J. Walker, Ed, A. Rodriguez, and J. Xiao, "Multi-view Self-supervised Deep Learning for 6D Pose Estimation in the Amazon Picking Challenge," *arXiv Computer Vision and Pattern Recognition*, no. 1609.09475, pp. 1–8, 2016. [Online]. Available: <http://apc.cs.princeton.edu/>
- [13] C. Hernandez, M. Bharatheesha, W. Ko, H. Gaiser, J. Tan, K. van Deurzen, M. de Vries, B. Van Mil, J. van Egmond, R. Burger, M. Morariu, J. Ju, X. Germann, R. Ensing, J. Van Frankenhuyzen, and M. Wisse, "Team Delft's Robot Winner of the Amazon Picking Challenge 2016," *arXiv Robotics*, no. 1610.05514, pp. 1–13, 2016.
- [14] J. Bohren and S. Cousins, "The SMACH High-Level Executive," *IEEE Robotics and Automation Magazine*, vol. 17, no. 4, pp. 18–20, 2010.
- [15] F. von Drigalski, L. El Hafi, P. M. Uriguen Eljuri, G. A. Garcia Ricardez, J. Takamatsu, and T. Ogasawara, "NAIST Drillbot: Drilling Robot at the Airbus Shopfloor Challenge," in *Proceedings of 34th Annual Conference of the Robotics Society of Japan (RSJ)*, 2016.
- [16] P. R. Wurman and J. M. Romano, "The Amazon Picking Challenge 2015," *IEEE Robotics and Automation Magazine*, vol. 22, no. 3, pp. 10–12, 2015.