Robots on Demand: A Democratized Robotics Research Cloud

Anonymous Author(s)
Affiliation
Address
`email`

Abstract: Robotics research is slowed by three challenges: building a robotics lab is expensive (few participants), everyone uses different robots (participants' findings often don't generalize outside their lab), and there is no internet-scale robotics dataset (no lab has the resources to make many robots do many different tasks to generate data and there is no data in the wild). The solution is to build a "Robotics Research Cloud" consisting of centers filled with remotely operable robots in standardized environments. This would be a valuable resource in pushing forward robot learning as a field by making cutting-edge robotics research broadly accessible, helping the field identify promising new approaches that succeed on agreed benchmarks, and creating a massive real-world robotics dataset similar to those that have revolutionized machine learning for vision and language.

Keywords: Remote robotics, open-source, benchmarking

1 Introduction

Robotics is stuck in its pre-ImageNet phase. Two major image classification competitions, the PASCAL VOC Challenge [1] and its much larger successor, the ImageNet Large Scale Visual Recognition Challenge [2], ushered in the modern era of machine learning. This was not only because of these datasets' scale but because they gave every aspiring machine learning researcher a platform on which to succeed: if you could build a program that empirically outperformed every other superstar researcher's attempt, then the field immediately paid attention to you. By expanding the number of participants in machine learning and defining a way by which they could agree on the best algorithms, ImageNet launched the modern machine learning revolution.

This transformation has not yet happened in robotics, for essentially the same reasons that plagued machine vision before ImageNet. The startup costs to creating a new robotics lab are prohibitively high: new investigators must invest not only hundreds of thousands of dollars for new robots and lab space, but also years of cumulative research hours setting up a new robot, calibrating it, and reimplementing and re-tuning numerous existing baselines just to begin contributing to the research frontier. Similarly, there is little guarantee that an algorithm achieving impressive results in one lab's setting will work well in others', due to myriad differences in robotic hardware, sensors, visual and physical properties of the testing environment, and other implementation details. As a result, even potential breakthroughs struggle to gain broader traction. Individual labs maintain their own understanding about which approaches work and build mostly on their own lab's previous work, limiting the benefits that reach the entire robotics community.

2 Related Efforts

Many roboticists have documented these challenges and have over the years tried to address them in different ways. Below is an overview of several such initiatives.

Standardizing hardware Robot platforms such as Willow Garage's PR2 and Rethink Robotics' Baxter attempt to standardize the hardware used across labs. Recent hardware efforts [3, 4, 5, 6, 7]

Submitted to the 5th Conference on Robot Learning (CoRL 2021). Do not distribute.

tend to involve exotic hardware targeting more niche communities. The YCB Object and Model set [8] standardizes the objects used in robotic environments. Standardized hardware is a useful step, but each robot is still only affordable to the biggest and best-funded labs, and environmental variations like lighting and sensing leave performance comparisons across labs difficult.

Standardizing software ROS [9] and follow-up efforts such as PyRobot [10] provide a common software stack allowing abstraction of parts of the robot hierarchy from perception to control. However, they do not address the benchmarking or access challenges described above.

Collecting and combining robotics data at scale Self-supervised experimentation (robots autonomously running experiments and evaluating their own success) allows individual labs to automatically collect large robotics datasets [11]. Google's "arm farm" further scaled data collection by using 14 robot arms working in parallel, substantially improving robotic grasping [12]. RoboNet [13], an open-source robotics dataset, facilitates data sharing across labs and has increased the scale of available robotics data. However, a dataset alone does not provide a way to easily evaluate models trained on the data or a way to compare algorithms across labs.

Simulation benchmarking Many simulators provide benchmarks that standardize results in both robotics settings [14, 15, 16, 17] and embodied AI settings with navigation components [18, 19, 20]. These have the benefit of eliminating variance due to physical environment conditions across lab setups. However, real-world performance is what we care about, where sample efficiency is more important and complex methods may be impossible to tune.

Simulation to reality Learning in simulation and transferring to reality is another promising approach [21, 22]. Performing most computation in simulation is an attractive way to scale robot learning, as simulation is safer and more efficient than the real world. OpenAI has shown success with sim2real, including dexterous manipulation of a Rubik's Cube [23]. However, it is next to impossible for a simulator to match the complexity and richness of the real world, making transfer inefficient. Many aspects of the real world, such as tactile sensing and dynamic interactions, are difficult to model, and learning in simulation is unlikely to yield real-world results in these domains.

In-person robotics competitions Some efforts take a different approach: running in-person robotics competitions. Entrants to these challenges run their own experiments and iterate independently before all teams congregate and finally test their methods on the real environment. Notable competitions include the DARPA challenges [24, 25, 26] and the Amazon Picking Challenge [27].

Remotely operable robotics testbeds An exciting new trend involves remotely operable robots, on which an experimenter can deploy code from anywhere in the world and observe the results using the environment's sensors. Duckietown [28] hosts the AI Driving Olympics, a competition on a series of tasks in a simplified autonomous driving world on low-cost standardized RC cars. Participants submit software solutions remotely which are run on physical robots in the environment. As an alternative model, the Georgia Tech Robotarium [29] allows anyone to remotely access a physical robotic swarm testbed, free for academic purposes. Their swarm consists of 20 low-cost RC cars in a shared space. The Robotarium also includes a simulator for code and safety checks before physical deployment. Finally, and most similarly to our proposal, the Real Robot Challenge [30] is an annual robotic manipulation competition allowing participants to test and compare their methods remotely on real hardware (seven TriFinger robots [5]) and a corresponding simulator.

3 Proposed Robotics Research Cloud

These recent efforts illuminate a path forward for robotics: remotely accessible robots on which everyone can run experiments, collect data, and benchmark their algorithms' performance. On top of this, collecting and anonymously releasing the recorded trajectories would create an ever-growing corpus of open-access robot operation data, unlocking large-scale machine learning applications in the robotics realm. All that's left is to put these ingredients together: a facility full of copies of the same robot set in standardized environments, connected to the internet for all researchers to access, fundamentally accelerating robotics as a field. Below, we sketch out a proposed structure.

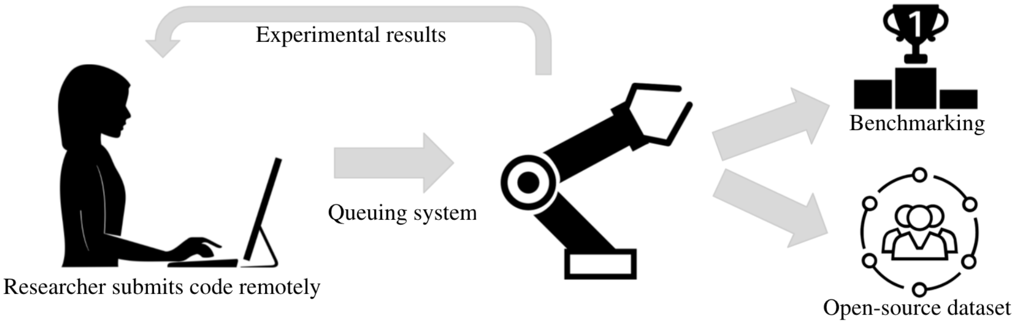

Figure 1: **Proposed workflow.** The proposed Robotics Research Cloud allows a researcher to quickly test an idea on a physical robot, obtain results quickly, and contribute to the community through benchmarks and open-sourced data.

3.1 Example Workflow

A prototypical use-case is visualized in Figure 1. A researcher might start with an existing codebase, make their own tweaks, define how long they'd like this code to run for and which benchmark task it's trying to satisfy (if any), and add it to a queue. The center prioritizes and runs these experiments on real robots (after they've passed simulation checks), scores them, and uploads the sensor and operation data to the cloud where the user can access it. Depending on the amount of training time needed, the experiment turnaround time might be 24 hours. On a slower timescale, anonymized trajectories would be added to an open-source dataset for community benefit.

3.2 Organizational Structure

Operation of the facility should be divided into a steering committee to guide policy decisions and ensure the center is being optimized for the many remote stakeholders and a local team, to handle the center's physical and digital infrastructure.

3.2.1 Steering Committee

The steering committee will be composed of robotics researchers and other stakeholders, responsible for decisions regarding the center's research agenda, high-level policies, and long-term development. These responsibilities would include:

- Choosing which robots and sensors to procure.
- Signing off on proposed new experiment environments and modification to existing environments based on researcher feedback.
- Determining access policies for the robots, including which institutions and researchers can run experiments and how their time is prioritized in the research queue.
- Determining the future trajectory of the Robotics Research Cloud, including whether and how to build additional centers.

3.2.2 Local Team

The local team will be responsible for building and maintaining the center itself. In addition to the handful of staff necessary to procure and maintain the robots, the local team may include dedicated staff necessary to reset the robots, in the experiments where environmental resets need human intervention. Importantly, the local team should include a team of full-time software engineers focused on building tooling to improve the experience of the remote research community. These tools should address the following challenges:

- Creating a unified software stack for robotic control, including a streamlined experience for remote researchers to schedule experiments, deploy code, and collect results. The codebase should allow researchers to easily take components from other research projects.
- Enable remote researchers to set up new benchmark tasks and create online leaderboards.
- Vet all code submitted to the facility against malicious content, including by maintaining a robot simulator of the robotic environment where checks must pass to ensure safety.

- Continuously publishing an anonymized dataset of robot trajectories collected in the lab.
- Maintaining remote community infrastructure, including virtual discussion spaces, workshops, and tutorials.

3.3 Center setup

The experimental specifications of the first Robotics Research Cloud center should be determined through discussion and consensus among robotics researchers in the field. One possible instantiation would be a manipulation focus, as manipulation tasks are more challenging than grasping alone but more feasible than navigation or locomotion tasks, which might require more space, safety checks, and manual environment resets. Additional focus areas could be added in subsequent centers. A manipulation center could include 100 Franka robot arms, each equipped with cameras, depth and tactile sensors. These robots could have environments that enable a few benchmark tasks, such as scooping, pouring, and writing. Initial methods attempting these benchmarks could be open-sourced as out-of-the-box baselines.

4 Discussion

4.1 Open Questions

Many open questions remain, the answers to which will impact the success of a Robotics Research Cloud.

The cloud will be most impactful if it achieves widespread adoption. Adoption depends on two key questions: Can researchers successfully prototype and run experiments remotely? Will researchers have the activation energy to adopt this new framework, when many labs already have their own? With high-quality infrastructure and a straightforward researcher interface, adoption is possible.

Choice of hardware and tasks will greatly affect research outcomes. Which sensors are necessary? How will sensor calibration and degradation be handled? Which tasks are feasible without relying heavily on manual environment resets? Through experimentation and iteration, we can find answers to these questions that will lead to sustainable operation of the Robotics Research Cloud.

With physical robots, safety is of paramount importance. While we can use simulation to run safety checks, how can we use an imperfect simulator to ensure safety in the real world?

Allocation of resources can be a contentious issue. How will queuing of experiments be prioritized? Should researchers be able to pay a fee to gain priority in the queue?

These questions should be discussed in detail by the community to ensure the success of the Robotics Research Cloud.

4.2 Limitations

A remote center does impose certain constraints, some of which we list here. Using fixed robots and environments rules out the possibility of jointly optimizing hardware and software. Centralizing robots also means less environment diversity as opposed to having robots individually acting in the wild. Remote teleoperation is high-latency, making collecting demonstrations or testing environment interaction difficult. Finally, latency in experimental results could increase iteration time.

4.3 Conclusion

By creating a facility for roboticists everywhere to run experiments and directly compare their results, we can give robotics its ImageNet breakthrough. A Robotics Research Cloud is likely to substantially accelerate the development of robotic software relative to its current trajectory. By massively increasing access to researchers across the field, it will reduce the network effects that concentrate talented robotics researchers in a small handful of schools and regions.

There are several next steps, including securing funding and convening robotics researchers to identify the initial experimental setup and benchmarks. Making the Robotics Research Cloud a reality will usher in a new age of remote robotics research and highly capable personal robots.

References

[1] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.

[2] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[3] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar. Robel: Robotics benchmarks for learning with low-cost robots. In *Conference on Robot Learning*, pages 1300–1313. PMLR, 2020.

[4] S. Cruciani, B. Sundaralingam, K. Hang, V. Kumar, T. Hermans, and D. Kragic. Benchmarking in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(2):588–595, 2020.

[5] M. Wüthrich, F. Widmaier, F. Grimminger, J. Akpo, S. Joshi, V. Agrawal, B. Hammoud, M. Khadiv, M. Bogdanovic, V. Berenz, et al. Trifinger: An open-source robot for learning dexterity. *arXiv preprint arXiv:2008.03596*, 2020.

[6] A. Gupta, A. Murali, D. Gandhi, and L. Pinto. Robot learning in homes: Improving generalization and reducing dataset bias. *arXiv preprint arXiv:1807.07049*, 2018.

[7] B. Yang, J. Zhang, V. Pong, S. Levine, and D. Jayaraman. Replab: A reproducible low-cost arm benchmark platform for robotic learning. *arXiv preprint arXiv:1905.07447*, 2019.

[8] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar. The ycb object and model set: Towards common benchmarks for manipulation research. In *2015 international conference on advanced robotics (ICAR)*, pages 510–517. IEEE, 2015.

[9] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng, et al. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.

[10] A. Murali, T. Chen, K. V. Alwala, D. Gandhi, L. Pinto, S. Gupta, and A. Gupta. Pyrobot: An open-source robotics framework for research and benchmarking. *arXiv preprint arXiv:1906.08236*, 2019.

[11] L. Pinto and A. Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 3406–3413. IEEE, 2016.

[12] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4-5):421–436, 2018.

[13] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn. Robonet: Large-scale multi-robot learning. *arXiv preprint arXiv:1910.11215*, 2019.

[14] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.

[15] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, pages 1094–1100. PMLR, 2020.

[16] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

[17] Y. Zhu, J. Wong, A. Mandlekar, and R. Martín-Martín. robosuite: A modular simulation framework and benchmark for robot learning. *arXiv preprint arXiv:2009.12293*, 2020.

[18] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.

[19] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. *arXiv preprint arXiv:2106.14405*, 2021.

[20] B. Shen, F. Xia, C. Li, R. Martín-Martín, L. Fan, G. Wang, S. Buch, C. D'Arpino, S. Srivastava, L. P. Tchapmi, et al. igibson, a simulation environment for interactive tasks in large realistic scenes. *arXiv preprint arXiv:2012.02924*, 2020.

[21] F. Sadeghi and S. Levine. Cad2rl: Real single-image flight without a single real image. *arXiv preprint arXiv:1611.04201*, 2016.

[22] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8973–8979. IEEE, 2019.

[23] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, et al. Solving rubik's cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.

[24] M. Buehler, K. Iagnemma, and S. Singh. *The DARPA urban challenge: autonomous vehicles in city traffic*, volume 56. springer, 2009.

[25] E. Krotkov, D. Hackett, L. Jackel, M. Perschbacher, J. Pippine, J. Strauss, G. Pratt, and C. Orlowski. The darpa robotics challenge finals: Results and perspectives. *Journal of Field Robotics*, 34(2):229–240, 2017.

[26] G. Seetharaman, A. Lakhotia, and E. P. Blasch. Unmanned vehicles come of age: The darpa grand challenge. *Computer*, 39(12):26–29, 2006.

[27] 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*, 15(1):172–188, 2016.

[28] L. Paull, J. Tani, H. Ahn, J. Alonso-Mora, L. Carlone, M. Cap, Y. F. Chen, C. Choi, J. Dusek, Y. Fang, et al. Duckietown: an open, inexpensive and flexible platform for autonomy education and research. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1497–1504. IEEE, 2017.

[29] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt. The robotarium: A remotely accessible swarm robotics research testbed. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1699–1706. IEEE, 2017.

[30] N. Funk, C. Schaff, R. Madan, T. Yoneda, J. U. De Jesus, J. Watson, E. K. Gordon, F. Widmaier, S. Bauer, S. S. Srinivasa, et al. Benchmarking structured policies and policy optimization for real-world dexterous object manipulation. *arXiv preprint arXiv:2105.02087*, 2021.