

OpenBot-Fleet: A System for Collective Learning with Real Robots

Matthias Müller, Samarth Brahmhatt, Ankur Deka, Quentin Leboutet, David Hafner, and Vladlen Koltun

Intel Labs

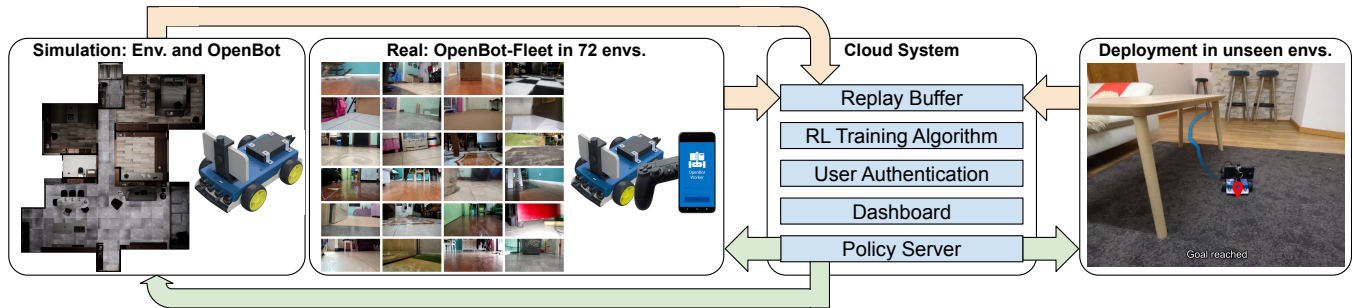


Fig. 1: Overview of OpenBot-Fleet, a comprehensive cloud robotics system for navigation which supports both simulated agents and real robots. Our wheeled robots are robust, low-cost, and use a smartphone for sensing and compute. A smartphone app allows for teleoperation and control by a navigation policy, and secure communication with the cloud system to upload robot experience and receive policies. We train a self-supervised perception backbone using data from 72 robots and pretrain a controller in simulation. The full navigation policy is evaluated in unseen real-world environments.

Abstract—We introduce **OpenBot-Fleet**, a comprehensive open-source cloud robotics system for navigation. OpenBot-Fleet uses smartphones for sensing, local compute and communication, Google Firebase for secure cloud storage and off-board compute, and a robust yet low-cost wheeled robot to act in real-world environments. The robots collect task data and upload it to the cloud where navigation policies can be learned either offline or online and can then be sent back to the robot fleet. In our experiments we distribute 72 robots to a crowd of workers who operate them in homes, and show that OpenBot-Fleet can learn robust navigation policies that generalize to unseen homes with $>80\%$ success rate. OpenBot-Fleet represents a significant step forward in cloud robotics, making it possible to deploy large continually learning robot fleets in a cost-effective and scalable manner. All materials can be found at <https://www.openbot.org/>.

I. INTRODUCTION

Learning high-performance policies for robot fleets deployed in diverse and changing environments is often difficult because of the mismatch between robots and environments used for training vs. deployment. Cloud robotics systems [1], [2], [3] can potentially address this issue by allowing deployed robots to collect and upload experience, improving the policy using this data, and sending the improved policy back to the robot fleet. However, while many open-source cloud robotics systems exist for different tasks [4], [5], [6], [7], almost none support the full policy improvement loop because of missing components like continuous policy deployment or experience data collection.

Offline learning (e.g., pretrained backbone networks for perception [8], [9]) can reduce the amount of online robot experience needed to train high-performance policies [10], [11]. However, training robotics backbone networks is challenging because of the expenses associated with owning robots and

operating them to act safely in real-world environments. Currently only large companies like Google and Tesla are able to operate large fleets of robots to acquire (private) datasets or implement the robot policy improvement loop at scale [12], [13], [14], [15].

In this paper, we focus on point-goal navigation [16] and present a comprehensive cloud robotics system that enables the full policy improvement loop and tackles the challenges of cost and scalability (see Figure 1). To the best of our knowledge, ours is the first open-source distributed robot navigation system that can collect large-scale experience in the real world and continuously improve and deploy a policy across a fleet of robots.

We adapt OpenBot [17] to develop a wheeled robot designed for mass production and ease of use by crowd workers in home environments (see Figure 2). Computation, sensing, and communication are done through a smartphone, leveraging the continuously improving capabilities of modern smartphones. We implement the full policy improvement software stack for both simulation (simulated environments, robot model, sensors) and the real world (cloud back end, smartphone app for robot control and sensor data logging) in a common learning framework. We demonstrate low-cost and scalable deployment by distributing our smartphone app and robot to a crowd of 72 workers, who provide the smartphone and home environment for learning point-goal navigation. The navigation policy consists of a perception module trained on data from this deployment and a control policy pretrained in simulation. This policy already achieves a success rate of 60% in unseen test environments compared to 15% by the best baseline. Online learning in these environments improves the success rate of our policy to 82.5%.

We make the following contributions: (1) A robust, low-cost wheeled robot with a companion smartphone app that supports robot control through teleoperation or a policy, user authentication, sensor data logging, and secure communication with the cloud back end. (2) A cloud back end with experience replay buffer, policy learning framework, and a dashboard. (3) System validation through experiments where we deploy a large fleet of robots to a crowd of workers to learn point-goal navigation policies that generalize to unseen environments and significantly outperform baselines.

II. RELATED WORK

Robot Learning Systems. Recent advances in robot learning, particularly real-world deep reinforcement learning (RL), have significantly improved robot performance in locomotion [18], [19], [20], navigation [21], [22], object manipulation [23], [12], [24], [10], and grasping [25], [26]. Despite the algorithmic advancements, the main challenges in robotics-oriented deep RL remain generalization and sample complexity [7], [18], [24]. To tackle these challenges, many researchers have focused on collecting large-scale datasets of multi-modal observations from multiple robots [25], [26], [27], [12], [7], [14], [28]. Additionally, the use of transformer models in combination with large-scale datasets has led to improved generalization performance [14], [28]. For example, RoboNet [7] and Gupta *et al.* [29] focus on vision-based manipulation. The former demonstrates that offline training of a perception backbone network with such a dataset enables finetuning for the target robot that is more sample efficient than robot-specific policy learning. The latter shows that diversity of the data collection conditions (environment, robot, etc.) is important for generalization. However, these systems are expensive and not fully open-source and reproducible (see Table I for details). In contrast, OpenBot-Fleet is affordable and completely open-source. In addition, while our experiments are designed for system validation, we observe similar effects with regard to sample efficiency and generalization.

Affordable Systems and Large-Scale Datasets. Researchers have introduced several systems for data collection, policy training and benchmarking [30], [31], [17], [32]. However, to the best of our knowledge, OpenBot-Fleet is the only system that implements and open-sources the full policy improvement loop: experience data logging from a distributed fleet, secure cloud storage, learning framework, and policy deployment to the fleet.

Point-goal Navigation. A prominent approach to solve this problem is SLAM [33] with optional visual landmark detectors and pose estimators [34], [35]. Recent map-free algorithms try to address challenges arising from dynamic environments and textureless scenes by learning through self-collision with the environment [36] or utilizing a transformer with the soft actor-critic algorithm [37]. Note that the focus of this paper is a comprehensive cloud robotics system capable of continual policy improvement rather than the development of a novel point-goal navigation algorithm.

III. SYSTEM

We summarize our system below and provide a more detailed description of each component in the supplement.

A. Robot Fleet

Each worker is given a wheeled robot, a remote control, and account credentials for logging into the smartphone app.

Robot. The robot is based on the DIY OpenBot [17] with several improvements vital for large-scale deployment while maintaining the low cost. The DIY OpenBot is 3D-printed and uses low-cost components like TT-motors with optical wheel encoders and an Arduino Nano. Building it is error-prone and time-consuming due to 3D-printing and manual tasks like wiring and soldering. To ensure durability and suitability for mass production, shipping, and repeated data collection with potential collisions, we have designed a robust injection mold plastic shell, fully integrated electronics with ATmega328 MCU, status LEDs and several sensors: wheel encoders, sonar, bump, and battery voltage (see Figure 2). A smartphone with our app for cloud robotics can easily be connected to the robot via a USB-C cable. Thus, our robot design addresses the challenges of cost and scalability encountered in cloud robotics.




Fig. 2: Our robot with a connected smartphone, adapted from OpenBot [17] for shipping to and use by crowd workers. In contrast to the original DIY version of OpenBot, it has a robust injection-molded plastic shell, integrated electronics (visible in the bottom-right image), and improved cable routing.

Smartphone App. Our smartphone app allows robot control through game controller commands or navigation policies after user login. It also logs sensor data for each episode and securely communicates experience data and policy weights to and from the cloud back end. Figure 3 provides an overview of all app features. Detailed descriptions, an app-flow diagram, and screenshots can be found in the supplement.

B. Policy Learning

Our system supports both offline learning, based on a dataset collected with the robot fleet, and online learning where a policy is continuously updated and redeployed based on continuously collected experience. In our method, we

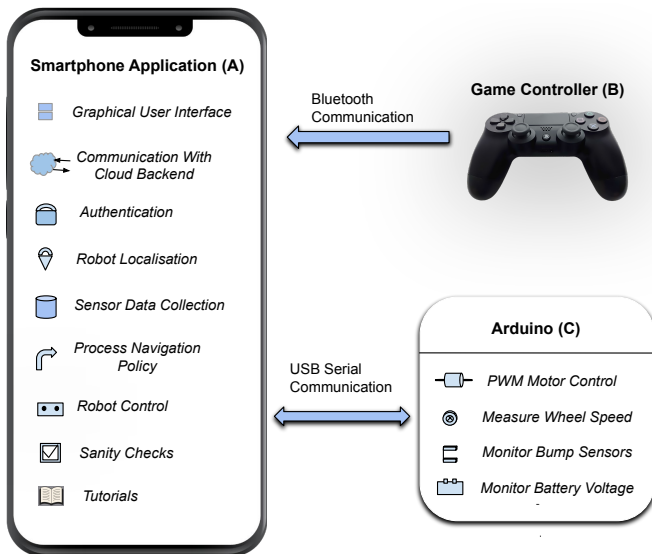


Fig. 3: Features of our app which runs on the smartphone connected to each robot. It communicates with the robot via a two-way serial link to send controls to it (e.g., while executing a control policy) and to receive sensor data for logging. A Bluetooth game controller can be used for teleoperation.

exploit both capabilities: (1) We train a perception backbone with data collected by crowd workers using self-supervised learning. (2) We update a control policy using reinforcement learning with real-world experience.

Offline Learning. Our framework supports training perception backbones on data collected with the fleet. We have implemented a lightweight but effective perception backbone based on U-Net [38]. Specifically, we train a floor segmentation network on the data collected by the robot fleet in a self-supervised fashion. We then sample points along the boundary as input representation for a control policy pretrained in simulation. This efficient and abstract representation allows for good sim-to-real transfer [39]. We have also implemented a pipeline for training end-to-end navigation policies using imitation learning. Please refer to Section IV-D (Baselines) for more details.

Online Learning. Our system also supports online learning with a fleet of robots deployed in the real world. This requires a cloud system for collecting experience from the robot fleet, updating the policy, and pushing it back to the fleet. To this end, we use Google Firebase [40] secure cloud storage.

Our learning framework is based on TensorFlow Agents [41] and the architecture is shown in Fig. 1-C of the supplement. The *driver* uses the current *collect policy* to execute an episode with the agent in the environment, and feeds the output *trajectory* to the *replay buffer*. Each *trajectory* contains a tuple of consecutive observations, action, reward, and additional meta-data, e.g., if a step caused the end of an episode. Note that the data collection and learning pipeline are asynchronous, so multiple agents can collect experience in different environments at the same time.

We implement a simulation driver (see Section III-C)

and a real-world driver, both of which can interact with the common replay buffer and policy update framework. The real-world driver leverages TensorFlow Lite [42] to run our TensorFlow policies efficiently on smartphones. While executing an episode, it records a log file containing the observations, policy action, and additional information like the robot pose relative to the target, for each step. Once the target is set by the worker as described in Section IV-A, we use ARCore [43], [44] to track the robot pose. After the episode is over, this log is uploaded to the cloud storage. The learning framework automatically retrieves it, converts it into the *trajectory* format, and adds it to the replay buffer. This design decision of not computing the reward directly on the smartphone allows greater flexibility in experimenting with the reward function later.

C. Simulation Platforms

We have integrated two simulators into our system: ProceduralSim, a fast procedural simulator, and SpearSim [45], a simulator based on UnrealEngine with photo-realistic indoor environments. We use ProceduralSim for pretraining a control policy using reinforcement learning. This learned controller can then be deployed on a real robot in conjunction with a perception backbone (see Section IV-A of the supplement) and can be further finetuned with our online learning framework. We use SpearSim mainly for evaluating our complete navigation stack before real-world deployment. Please refer to the supplement for more details.

D. Comparison to Other Systems

Table I presents a comparison of our proposed distributed learning system to other works. While the largest deployed fleets of robots in industry, i.e., autonomous cars and robots for warehouse logistics and last mile delivery, require good navigation policies, most existing research systems for large-scale learning focus on manipulation. In contrast, our system focuses on navigation. The small cost compared to other systems also enables much larger deployments. Our system is designed to be deployed in the wild and the user interface through a smartphone app does not require robotics experts for operation. Other systems are usually deployed in controlled lab settings and require expert supervision. To reproduce a system, code and instructions for re-creating the system (including the robots) need to be available. Similarly, data, models and code need to be released for model reproducibility. While other works only release one or the other, we open-source and release all of it.

IV. METHODOLOGY

We implement a robotic navigation system that can drive a robot from A to B while avoiding obstacles.

A. Large-Scale Data Collection

We collect large-scale driving data with the help of crowd workers. The task is to drive from A to B, taking the shortest path while avoiding obstacles. The data collected consists of image-control pairs and target locations. We use this data

System	Robots			Envs		Tasks		Reproducible	
	Type	Cost [\$]	#	Type	#	Man.	Nav.	Sys.	Mod.
Gupta <i>et al.</i> [29]	MM	3,000	5	W	6	✓	✗	✓	✗
RoboNet [7]	SM	26,000	7	L	4	✓	✗	✗	✓
REPLAB [31]	SM	3,000	2	L	3	✓	✗	✓	✗
QT-Opt [12]	SM	25,000	7	L	7	✓	✗	✗	✗
RT-1 [14]	MM	N/A*	13	L	≥8	✓	✓	✗	✓
Ours	GV	100	72	W	72	✗	✓	✓	✓

TABLE I: Comparison of our proposed multi-robot distributed learning system to other similar works. Robot types are denoted by MM (Mobile Manipulator), SM (Static Manipulator), and GV (Ground Vehicle). Environment types are denoted by W (in-the-wild deployment) and L (controlled lab setting). Manipulation, Navigation, System, and Model are abbreviated with the first three letters. *Robots not available for purchase.

for (1) learning a conditional imitation learning [46] baseline which regresses controls end-to-end from images and relative target poses, and (2) training our self-supervised perception backbone from images.

The data collection process involved 72 crowd workers using Bluetooth controllers and smartphones attached to the robot bodies to record data in diverse natural indoor home environments such as living rooms, bedrooms, and kitchens. The data contains images collected by the smartphone camera, synchronized with controls supplied by the human operators and state estimation computed with ARCore. The complete dataset consists of 15,642 recordings and over 1M frames with a variable frame rate of up to 10 frames per second depending on the smartphone capability.

Participants were given specific instructions to ensure the quality of the data collected. They were instructed to place the robot on the ground before commencing the recording, to drive smoothly from the starting point to the goal without stopping, using the shortest path while maneuvering around obstacles such as furniture, walls, and naturally placed household items. Also, to ensure that the trajectory length is at least 2 meters and to prevent the robot from being picked up during the recording, an automatic check was put in place. To ensure a diverse set of data, participants were instructed to change the start and goal locations at least after every 5 to 10 episodes. They were also instructed to avoid certain behaviors during the episodes such as driving in circles or zig-zagging, and stopping the robot during the task unless they had reached the end point or the robot had collided. To ensure privacy, participants were also instructed to only record in their own home environment, avoid people’s faces and actively review each recording before submitting it. Despite best efforts, the data quality varies based on worker experience and motivation, robot, and environment condition.

B. Learning Perception

Our collected dataset is large, diverse, and suitable for training and evaluating various perception modules for the robot. We train two modules. (1) A end-to-end baseline

(see Section IV-D) that regresses the action to be applied for the next step from a given RGB image supervised by the synchronized teleoperation command. (2) For our system, a perception backbone that predicts the floor boundary from a given RGB image.

The perception backbone for the floor boundary prediction is trained in a scalable self-supervised manner which does not require human labelling. We generate an initial floor mask by applying a gradient threshold on the depth image generated by MiDaS [47] from the RGB image and further refine the mask using GrabCut [48]. For more details, see Section IV-A in the supplement. We use the more than 1M images collected by the crowd workers and an additional 12k images from SpearSim (to reduce the sim-to-real gap between evaluation in simulation and the real world). The resulting dataset contains 1,068,577 images in total and is split randomly into 90% for training and 10% for validation. The network is based on UNet [38] and trained to segment the floor. The boundary is the uppermost pixel location of this mask for each image column, and is used as an observation representation for our navigation policy.

C. Learning Control

We train a control policy using reinforcement learning. The policy is a residual network with respect to a unicycle model. We first pretrain in simulation using ground truth observations to collect a large amount of experience quickly. We then finetune the policy using the real robot fleet.

Unicycle Model. By default, the robot moves towards the goal using a unicycle model:

$$v_l = \frac{\cos \alpha - \sin \alpha}{\sqrt{2}} \quad \text{and} \quad v_r = \frac{\cos \alpha + \sin \alpha}{\sqrt{2}}, \quad (1)$$

where v_l, v_r denote the left and right motor commands, and α is the target heading angle w.r.t. the robot front. The learnt policy predicts residual actions $\Delta v_l, \Delta v_r$ which are added to v_l, v_r to get the final motor commands.

Learning Algorithm. We use the popular off-policy RL algorithm Soft Actor-Critic (SAC) [49]. Thanks to the use of a large replay buffer where past experiences are stored and re-used to train the policy, SAC is usually more sample efficient than on-policy RL algorithms like Proximal Policy Optimization (PPO) [50]. Additionally, SAC has an entropy bonus that encourages exploration.

Training Process. Once a training session is started, the perception backbone, control policy and training meta-data are uploaded to our cloud back end. The worker’s app gets notified and the worker can collect an episode recording and upload it. The uploaded recording is stored in the cloud and immediately passed on to the training session, where we add it to the replay buffer and update the RL policy. The updated policy is published to the workers and the cycle repeats.

D. Real-world deployment and evaluation

We deploy our policy to real robots in different unseen environments (see Figure 4) to evaluate both the zero-shot

and few-shot performance. We compare our approach to several baselines.



Fig. 4: Unseen real-world environments used for evaluation.

Task Definition. We evaluate our system on the standard and yet challenging task of point-goal navigation [16]. Here, the goal is to learn a navigation policy that automatically drives the robot (without collisions) to a user-defined target in an indoor environment. We have chosen point-goal navigation as the optimal task for evaluating our system due to several reasons. Firstly, most smartphones support ARCore which can be used for state estimation, defining the goal position, and computing the reward for the reinforcement learning algorithm. Secondly, the point-goal navigation task is simple to define, yet a challenging problem to solve. This task is also useful for practical applications in real-world environments, *e.g.*, autonomous vacuum cleaners or home assistants. Finally, this task is well-suited for evaluating the effectiveness of our proposed approach in solving navigation problems at scale.

Evaluation Protocol and Metrics. We regard a policy rollout as successful if the robot reaches the target in a specified time without collision. A collision is either automatically detected by the bump sensors or identified by user intervention. We apply three metrics to evaluate a policy: success rate, distance to the goal at the final time step, and collision rate. An episode ends when either the goal is reached, the robot collides, or the maximum number of steps is reached (timeout). On the real robot, episodes can also end due to ARCore failures.

Baselines. We implement a goal-conditioned imitation learning network [46] with PilotNet backbone [51] that regresses controls end-to-end based on ARCore goal and images (CIL [46] - PilotNet [51]). For this baseline, the data collected by the crowd workers is first split into a training set and a validation set, with 65 workers contributing 14,141 recordings and 957,434 frames to the training set, and 7 workers contributing 1,501 recordings and 98,198 frames to the validation set. The data is then filtered to remove bad recordings and frames where no control signal was supplied (*e.g.*, robot is stationary) resulting in 590,465 image-control pairs in the training set and 60,291 image-control pairs in the validation set. This baseline provides a comparison for the performance of our proposed method to behaviour cloning.

We also implement two other baselines, representing different observation representations: (1) *MobileNetV2*: image features extracted with MobileNetV2 [52] trained on ImageNet [8], and (2) *DepthRays*: the depth values of a pencil of 11 rays parallel to the ground and originating at the camera center, normalized in the $[0, 1]$ range. Both representations are used to train a navigation policy with reinforcement learning in simulation in the same fashion as for our approach. The simulation training phase uses the ground truth RGB rendering and depth image respectively, similar to the simulation training phase of our policy, which uses the ground truth floor boundary. The real world deployment and evaluation phase uses the real RGB image and the real depth image generated by ARCore [44], [43].

V. RESULTS

A. Zero-shot generalization

Our evaluation on real robots in four unseen environments is shown in Table II. Our method achieves superior zero-shot performance with a success rate (SR) of at least 50% and up to 80%. The next best method is DepthRays, achieving 30% in two of the environments. However, it completely fails in the other two (0% SR). This is probably due to the noisy and unreliable depth compared to simulation. MobileNetV2 performs even worse with 0% SR in all environments but one. We conjecture that this is due to the distribution mismatch – the camera mounted on the robot is very low to the ground and images contain mostly floor and furniture whereas most images in ImageNet are captured by humans and focus on diverse objects in various settings. The behaviour cloning baseline, CIL-PilotNet, also performs poorly despite the additional supervision from the control labels by the crowd workers. Behaviour cloning is sensitive to bad samples [53], and not all of the workers can be considered experts. In addition, they may have different driving styles and optimize for task completion rather than trajectory quality.

B. Few-shot adaptation

We also investigate the potential of our method to improve performance with online adaptation in the real world. To this end, we deploy our policy in different unseen environments and finetune the policy on up to 200 episodes using our online learning system. Our method achieves at least 80% and up to 90% success rate across all environments with less than 200 episodes of online learning. Note that state estimation by ARCore is noisy which can lead to unreachable goals making a success rate of 100% challenging to impossible.

We also experimentally confirm that the finetuned policies do not overfit to the environments. For example, the policy finetuned in environment 1, still achieves 70% success rate in environment 4. We also compare the adaptation performance of our method to the baselines with up to 4,000 episodes of finetuning per policy using the photo-realistic simulator SpearSim. In summary, only our method is able to quickly adapt and achieve high performance.

	Environment 1			Environment 2			Environment 3			Environment 4			Average		
	SR \uparrow	GD \downarrow	CR \downarrow	SR \uparrow	GD \downarrow	CR \downarrow	SR \uparrow	GD \downarrow	CR \downarrow	SR \uparrow	GD \downarrow	CR \downarrow	SR \uparrow	GD \downarrow	CR \downarrow
CIL [46] - PilotNet [51]	0	1.07	100	10	0.49	80	0	1.86	100	0	1.68	100	2.5	1.28	95.0
MobileNetV2 [52] (zero-shot)	0	0.27	70	10	1.14	40	0	2.73	70	10	2.07	90	5.0	1.55	67.5
DepthRays [43] (zero-shot)	30	0.80	70	30	0.81	60	0	1.87	100	0	2.19	100	15.0	1.42	82.5
Ours (zero-shot)	<u>50</u>	<u>0.31</u>	40	<u>60</u>	<u>0.51</u>	<u>30</u>	<u>50</u>	<u>0.79</u>	<u>50</u>	<u>80</u>	<u>0.09</u>	<u>20</u>	<u>60.0</u>	<u>0.43</u>	<u>35.0</u>
Ours (≤ 200 episodes)	80	0.20	20	80	0.02	10	80	0.27	20	90	0.02	10	82.5	0.13	15.0

TABLE II: Point-goal navigation results with real robot in unseen environments. *SR*: success rate [%], *GD*: goal distance at the end of the episode [m], *CR*: collision rate [%]. **Bold** for best, underlined for second-best. Models are pretrained for 20k episodes in ProceduralSim and use a history length of 5. All models are evaluated for 10 trials and results are averaged. We finetune models until they achieve a success rate of 100% in 10 consecutive training episodes or reach 200 episodes.

# of envs.	mIOU \uparrow	MAE \downarrow	SR \uparrow	GD \downarrow	CR \downarrow
1	90.95	4.58	50	0.69	50
6	91.94	4.01	70	0.44	30
72	93.49	2.95	85	0.095	15

TABLE III: Ablation study of the perception backbone. We train with data from 1, 6, and 72 environments and show that diversity and scale improve both the accuracy of the perception backbone and the performance of the navigation policy. *mIOU*: mean intersection over union for floor segmentation [%], *MAE*: mean average error for floor boundary, *SR*: success rate [%], *GD*: goal distance at the end of the episode [m], *CR*: collision rate [%].

C. Ablation studies

To study the impact of diversity and scale, we train the perception backbone on data from 1, 6, and 72 workers. In the case of a single environment, we pick one that is relatively similar to the test environment used for backbone evaluation and sample 217 recordings (the average number of recordings per environment) for training a perception backbone. Similarly, for the case of 6 environments we train a model on a total of 1302 recordings (6×217) across 6 environments. Finally, we also train a model on all data from all 72 environments. We then evaluate the performance of the perception backbone on our test set and also evaluate the zero-shot task-performance in different novel environments.

The results presented in Table III show the correlation between perception backbone metrics and zero-shot task performance. We observe a clear trend: as the diversity and quantity of data increases, the performance metrics improve. Specifically, the mIOU increased from 90.95% for 1 environment to 93.49% for 72 environments, indicating the effectiveness of learning at scale and in diverse environments.

Furthermore, we evaluate the perception backbone in conjunction with our controller trained in simulation on two novel unseen real-world environments. The success rate (SR), collision rate (CR), and average remaining distance to the goal (GD) were measured across 20 trials. The results show that increasing the number of training environments leads to improved task performance. For instance, when trained with data from 72 environments, the system achieves a SR of 85%, a CR of 15%, and an average remaining

distance to the goal of 0.095m.

These findings highlight the importance of both diversity and scale in training good navigation policies and demonstrates the usefulness of our cloud robotics system. The correlation between the perception backbone metrics and task performance confirms the effectiveness of our crowd-sourced data collection approach. Collecting data from a larger number of diverse environments not only improves the perception backbone performance but also enhances the system’s overall performance in real-world scenarios. The presented results validate the effectiveness and impact of our system in achieving robust and reliable performance in cloud robotics applications. We provide further ablation studies for selecting the parameters of our approach in the supplement.

VI. CONCLUSION

We presented a comprehensive cloud robotics system that addresses cost and scalability challenges in learning navigation policies with robot fleets in real-world environments. It consists of a low-cost, plug-and-play robot equipped with a robust injection-molded plastic shell, integrated electronics, and a software stack that enables data collection and online learning for large fleets of robots. We demonstrated the effectiveness of our system through experiments with a fleet of 72 robots distributed to crowd workers. The resulting navigation policies exhibit good performance in unseen environments with less than 200 episodes of online learning.

While the results are promising, there are some limitations. The robot’s actuation is noisy, limiting zero-shot performance in new environments. ARCore, used for state estimation, exhibits drift and frequently loses tracking since it was not designed with robot navigation in mind. Beyond modeling the noisy robot actuation and improving state estimation, it would also be interesting to scale up to fleets beyond 1,000 and to continuously update both the perception backbone and the control policy.

A video with a short summary of this work can be found at <https://youtu.be/3tuKzbtDFcc>. The supplementary material which contains further implementation details is available at <https://t.ly/v9be9>.

REFERENCES

- [1] B. Kehoe, S. Patil, P. Abbeel, and K. Goldberg, "A survey of research on cloud robotics and automation," *IEEE Transactions on automation science and engineering*, vol. 12, no. 2, pp. 398–409, 2015.
- [2] J. Wan, S. Tang, H. Yan, D. Li, S. Wang, and A. V. Vasilakos, "Cloud robotics: Current status and open issues," *Ieee Access*, vol. 4, pp. 2797–2807, 2016.
- [3] O. Saha and P. Dasgupta, "A comprehensive survey of recent trends in cloud robotics architectures and applications," *Robotics*, vol. 7, no. 3, p. 47, 2018.
- [4] 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)*. IEEE, 2017, pp. 1497–1504.
- [5] 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)*, 2017, pp. 1699–1706.
- [6] S. Bauer, M. Wüthrich, F. Widmaier, A. Buchholz, S. Stark, A. Goyal, T. Steinbrenner, J. Akpo, S. Joshi, V. Berenz, V. Agrawal, N. Funk, J. Urain De Jesus, J. Peters, J. Watson, C. Chen, K. Srinivasan, J. Zhang, J. Zhang, M. Walter, R. Madan, T. Yoneda, D. Yarats, A. Allshire, E. Gordon, T. Bhattacharjee, S. Srinivasa, A. Garg, T. Maeda, H. Sikchi, J. Wang, Q. Yao, S. Yang, R. McCarthy, F. Sanchez, Q. Wang, D. Bulens, K. McGuinness, N. O'Connor, R. Stephen, and B. Schölkopf, "Real robot challenge: A robotics competition in the cloud," in *Proceedings of the NeurIPS 2021 Competitions and Demonstrations Track*, ser. Proceedings of Machine Learning Research, D. Kiela, M. Ciccone, and B. Caputo, Eds., vol. 176. PMLR, 06–14 Dec 2022, pp. 190–204. [Online]. Available: <https://proceedings.mlr.press/v176/bauer22a.html>
- [7] 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:1910.11215*, 2019.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [9] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, and J. Lee, "Transporter networks: Rearranging the visual world for robotic manipulation," *Conference on Robot Learning (CoRL)*, 2020.
- [10] A. Herzog, K. Rao, K. Hausman, Y. Lu, P. Wohlhart, M. Yan, J. Lin, M. G. Arenas, T. Xiao, D. Kappler, *et al.*, "Deep RL at Scale: Sorting Waste in Office Buildings with a Fleet of Mobile Manipulators," *arXiv preprint arXiv:2305.03270*, 2023.
- [11] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Proceedings of the 5th Conference on Robot Learning (CoRL)*, 2021.
- [12] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, *et al.*, "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation," *arXiv preprint arXiv:1806.10293*, 2018.
- [13] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, 2020.
- [14] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, *et al.*, "Rt-1: Robotics transformer for real-world control at scale," *arXiv preprint arXiv:2212.06817*, 2022.
- [15] S. Wu, B. Yu, S. Liu, and Y. Zhu, "Autonomy 2.0: The quest for economies of scale," *arXiv preprint arXiv:2307.03973*, 2023.
- [16] P. Anderson, A. X. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, and A. R. Zamir, "On evaluation of embodied navigation agents," *CoRR*, vol. abs/1807.06757, 2018. [Online]. Available: <http://arxiv.org/abs/1807.06757>
- [17] M. Müller and V. Koltun, "Openbot: Turning smartphones into robots," in *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2021.
- [18] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 698–721, 2021.
- [19] L. Smith, I. Kostrikov, and S. Levine, "A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning," *arXiv preprint arXiv:2208.07860*, 2022.
- [20] D. Gandhi, L. Pinto, and A. Gupta, "Learning to fly by crashing," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [21] S. Gattu, "Autonomous navigation and obstacle avoidance using self-guided and self-regularized actor-critic," in *Proceedings of the 8th International Conference on Robotics and Artificial Intelligence*, 2022, pp. 52–58.
- [22] S. Levine and D. Shah, "Learning robotic navigation from experience: principles, methods and recent results," *Philosophical Transactions of the Royal Society B*, vol. 378, no. 1869, p. 20210447, 2023.
- [23] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, "Manipulation by feel: Touch-based control with deep predictive models," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 818–824.
- [24] Y. Lu, K. Hausman, Y. Chebotar, M. Yan, E. Jang, A. Herzog, T. Xiao, A. Irpan, M. Khansari, D. Kalashnikov, and S. Levine, "AW-Opt: Learning Robotic Skills with Imitation and Reinforcement at Scale," in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, D. Hsu, and G. Neumann, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 1078–1088. [Online]. Available: <https://proceedings.mlr.press/v164/lu22a.html>
- [25] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *ICRA*, 2016.
- [26] 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," *I. J. Robotics Res.*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [27] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *Conference on Robot Learning (CoRL)*, 2018.
- [28] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, *et al.*, "A generalist agent," *arXiv preprint arXiv:2205.06175*, 2022.
- [29] A. Gupta, A. Murali, D. P. Gandhi, and L. Pinto, "Robot learning in homes: Improving generalization and reducing dataset bias," in *NeurIPS*, 2018.
- [30] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar, "Robel: Robotics benchmarks for learning with low-cost robots," *arXiv:1909.11639*, 2019.
- [31] B. Yang, J. Zhang, V. Pong, S. Levine, and D. Jayaraman, "Replab: A reproducible low-cost arm benchmark platform for robotic learning," *arXiv:1905.07447*, 2019.
- [32] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, *et al.*, "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *Science Robotics*, vol. 7, no. 67, p. eabl6259, 2022.
- [33] F. Blochliger, M. Fehr, M. Dymczyk, T. Schneider, and R. Siegwart, "Topomap: Topological mapping and navigation based on visual slam maps," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 3818–3825.
- [34] S. Frintrop and P. Jensfelt, "Attentional landmarks and active gaze control for visual slam," *IEEE Transactions on Robotics*, vol. 24, no. 5, pp. 1054–1065, 2008.
- [35] S. Brahmabhatt, H. B. Amor, and H. Christensen, "Occlusion-aware object localization, segmentation and pose estimation," in *Proceedings of the British Machine Vision Conference (BMVC)*. BMVA Press, September 2015, pp. 80.1–80.13. [Online]. Available: <https://dx.doi.org/10.5244/C.29.80>
- [36] G. Kahn, A. Villafior, B. Ding, P. Abbeel, and S. Levine, "Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 5129–5136.
- [37] W. Huang, Y. Zhou, X. He, and C. Lv, "Goal-guided transformer-enabled reinforcement learning for efficient autonomous navigation," *arXiv preprint arXiv:2301.00362*, 2023.
- [38] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.

- [39] M. Müller, A. Dosovitskiy, B. Ghanem, and V. Koltun, "Driving policy transfer via modularity and abstraction," *arXiv preprint arXiv:1804.09364*, 2018.
- [40] "Google firebase," <https://firebase.google.com>, accessed: 2023-09-15.
- [41] S. Guadarrama, A. Korattikara, O. Ramirez, P. Castro, E. Holly, S. Fishman, K. Wang, E. Gonina, N. Wu, E. Kokiopoulou, L. Sbaiz, J. Smith, G. Bartók, J. Berent, C. Harris, V. Vanhoucke, and E. Brevdo, "TF-Agents: A library for reinforcement learning in tensorflow," 2018, [Accessed 23-Aug-2021]. [Online]. Available: <https://github.com/tensorflow/agents>
- [42] TensorFlow Lite, "<https://www.tensorflow.org/lite>," accessed: 2023-09-15.
- [43] R. Du, E. Turner, M. Dzitsiuk, L. Prasso, I. Duarte, J. Dourgarian, J. Afonso, J. Pascoal, J. Gladstone, N. Cruces, *et al.*, "Depthlab: Real-time 3d interaction with depth maps for mobile augmented reality," in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, 2020, pp. 829–843.
- [44] ARCore, "<https://developers.google.com/ar>," accessed: 2023-09-15.
- [45] M. Roberts, Q. Leboutet, R. Prakash, R. Wang, H. Zhang, R. Tang, M. Ferragut, S. Leutenegger, S. R. Richter, V. Koltun, M. Müller, and G. Ros, "SPEAR: A simulator for photorealistic embodied ai research," <http://github.com/isl-org/spear>, 2022.
- [46] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *ICRA*, 2018.
- [47] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 3, pp. 1623–1637, 2020.
- [48] C. Rother, V. Kolmogorov, and A. Blake, "GrabCut – interactive foreground extraction using iterated graph cuts," *ACM transactions on graphics (TOG)*, vol. 23, no. 3, pp. 309–314, 2004.
- [49] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *ICML*, 2018.
- [50] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [51] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, "End to end learning for self-driving cars," *arXiv:1604.07316*, 2016.
- [52] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv:1704.04861*, 2017.
- [53] F. Sasaki and R. Yamashina, "Behavioral cloning from noisy demonstrations," in *International Conference on Learning Representations*, 2021.