

Scaling Robot Supervision to Hundreds of Hours with RoboTurk: Robotic Manipulation Dataset through Human Reasoning and Dexterity

Ajay Mandlekar¹, Jonathan Booher¹, Max Spero¹, Albert Tung¹, Anchit Gupta¹,
Yuke Zhu¹, Animesh Garg^{2,3}, Silvio Savarese¹, Li Fei-Fei¹

Abstract—Large, richly annotated datasets have accelerated progress in fields such as computer vision and natural language processing, but replicating these successes in robotics has been challenging. While prior data collection methodologies such as self-supervision have resulted in large datasets, the data can have poor signal-to-noise ratio. By contrast, previous efforts to collect task demonstrations with humans provide better quality data, but they cannot reach the same data magnitude. Furthermore, neither approach places guarantees on the diversity of the data collected, in terms of solution strategies. In this work, we leverage and extend the RoboTurk platform to scale up data collection for robotic manipulation using remote teleoperation. The primary motivation for our platform is two-fold: (1) to address the shortcomings of prior work and increase the total quantity of manipulation data collected through human supervision by an order of magnitude without sacrificing the quality of the data and (2) to collect data on challenging manipulation tasks across several operators and observe a diverse set of emergent behaviors and solutions. We collected over 111 hours of robot manipulation data across 54 users and 3 challenging manipulation tasks in 1 week, resulting in the largest robot dataset collected via remote teleoperation. We evaluate the quality of our platform, the diversity of demonstrations in our dataset, and the utility of our dataset via quantitative and qualitative analysis. For additional results, supplementary videos, and to download our dataset, visit roboturk.stanford.edu/realrobotdataset

I. INTRODUCTION

Crowdsourcing mechanisms such as Amazon Mechanical Turk have facilitated the creation of large, richly annotated datasets. The advent of datasets, sizing in millions, has accelerated progress in computer vision and natural language processing [7, 30] by enabling the development and evaluation of a wide range of learning algorithms and benchmarks. Efforts to aggregate similar amounts of data promise to boost performance in the field of robot manipulation.

Subsequently, the community leveraged online self-supervised data collection [22, 28] and off-policy reinforcement learning [16] to collect large quantities of physical robot data for tasks such as grasping (over 1000 hours). However, the data collected through such methods often has low signal-to-noise ratio, since a large portion of the data is collected by applying random controls. Subsequently, the time it takes to start collecting high quality data can be prohibitively large, limiting the complexity of the tasks achievable with this approach. Furthermore, specification and evaluation of a reward function for complex tasks can be non-intuitive. In contrast, human demonstrations obviate the

¹Stanford AI Lab (SAIL), Stanford University. ²Nvidia, USA. ³University of Toronto, Canada. Email: amandlek@stanford.edu

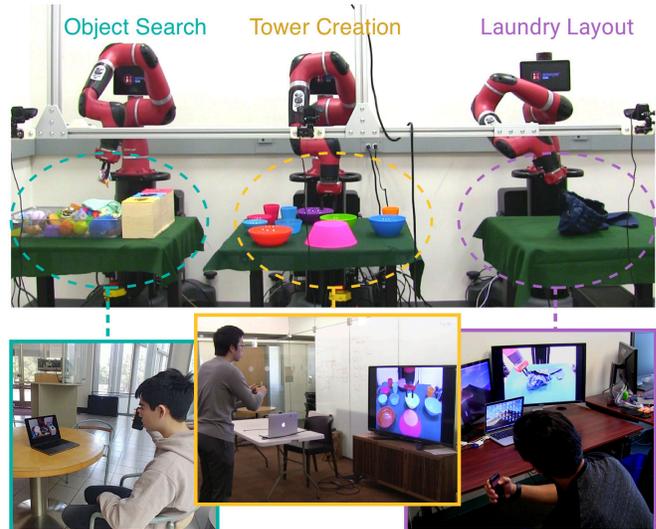


Fig. 1: Collecting data on physical robot arms with the RoboTurk platform. To collect task demonstrations, users connect to our platform from remote locations using a web browser and use their smartphone as a motion controller to move the physical robot arm in free space. Users are provided a video stream of the robot workspace in their web browser.

need for this specification by implicitly providing a set of successful task executions.

Prior work [8, 43] has shown that imitation learning on data collected from humans can achieve success on a set of restricted task instances [1, 3, 5, 21, 29, 31, 38]. However, these approaches have been limited in both the scale of data collected and the complexity of the tasks used for data collection. The ideal approach would be able to collect data on the scale of self-supervised methods but with the quality of human-in-the-loop approaches.

However, replicating the success and impact that large-scale datasets have had in vision and language for robotics has been challenging. The core problem is that the expert needs to *demonstrate* how to perform a task in real-time, instead of offline data-labeling. Therefore, methods for real-time remote interaction that are robust to delays in both actuation and network latency must be established. More importantly these methods must operate at scale to facilitate crowdsourcing.

This paper addresses the problem of large scale crowdsourcing on real robots. We propose an approach to collect task demonstrations on physical robots from humans to scale data collection along two dimensions: the quantity of data collected, and the diversity of data collected. We extend the RoboTurk platform [25] from simulation to real robots, and address the ensuing challenges such as: establishing a

framework to remotely manage robots, handling additional delays due to hardware actuation not present in simulators, ensuring operational safety when operated by novice users, and managing alignment of data streams from a multitude of sensors for a number of robots simultaneously. Furthermore, we present three robotic manipulation tasks that require human intervention both at the level of reasoning and dexterity of manipulation. Lastly, we present a large-scale dataset on these three tasks, comprised of over 111 hours of data collected by 54 people resulting in a diverse set of solutions that is an order of magnitude larger than state-of-the-art.

Summary of Contributions:

- 1) We extend the RoboTurk crowdsourcing platform to enable remote data collection on physical robots. Our new platform accounts for limited robot resources, additional delays introduced by robot hardware, the need for safety measures to protect the robots from harm, and data collection from multiple sensor streams at different rates.
- 2) We introduce Object Search, Tower Creation, and Laundry Layout: three different tasks that require human intervention in the form of low-level dexterity of manipulation and high-level cognitive reasoning to solve the tasks.
- 3) We present the largest robot dataset collected via remote teleoperation. Over the course of 1 week, we collected over *111 hours* of data across *54 users* on the 3 challenging manipulation tasks that we introduced.
- 4) We evaluate the quality of our platform, the diversity of demonstrations in our dataset, and the utility of our dataset via quantitative and qualitative analysis.

II. RELATED WORK

Large-Scale Data collection in Robotics. Data-driven methods for learning in robotics have been used to collect grasps [12] and object models [17], and run large scale physical trials for grasping [16, 22, 28] and pushing [41]. These methods used hundreds of hours of robot interaction, although a majority of the trials were not successful.

Simulated and Self-supervised Methods. Large scale self-supervision has low signal-to-noise ratio due to exploration via a random policy. While simulators can scale easily and provide many task variations, several task types, such as those shown in this work, can be difficult to simulate. Combinations of these methods as in [15, 24], are limited by simulator fidelity, and often focused on tasks with specific and easily measurable success criterion.

Learning from Demonstration and Imitation Learning. Imitation learning (IL) is often preferred over RL to achieve efficiency in policy learning. Specification of reward functions can be non-intuitive for a number of robotic tasks [27]. Imitation learning can be performed mainly through inverse reinforcement learning (IRL) [2, 21] or behavioral cloning (BC) [29, 31, 32]. However, these algorithms typically either require a large amount of data (BC) or a large number of environment interactions (IRL).

Crowdsourced Teleoperation for Robot Learning. Collecting large amounts of data has been a challenge for continuous

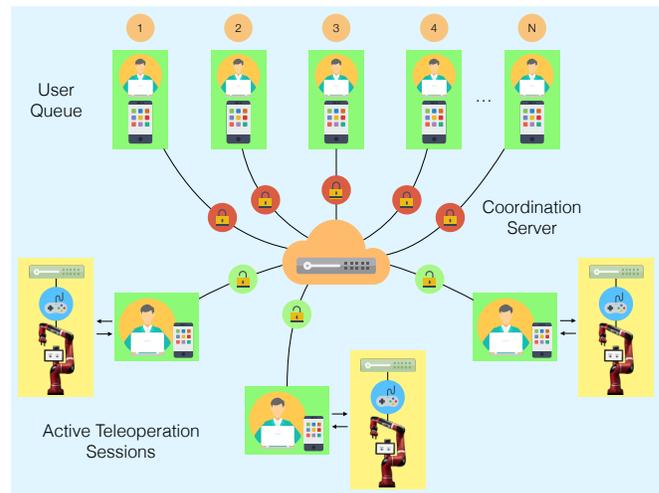


Fig. 2: **RoboTurk system diagram for data collection on physical robot arms.** The coordination server manages users and guards access to the physical robot arms. There is a limited number of active teleoperation sessions (one per robot arm). Other users are queued in the system until one of the robot arms is available for data collection. When this occurs, the coordination server creates a new connection between the robot arm and the user at the front of the queue, and a new teleoperation session begins.

manipulation tasks. Crowdsourcing supervision has resulted in some remarkable scaling of datasets in computer vision and natural language [7, 30]. In robotics, crowdsourcing over the internet was first introduced to robotics in the Telegarden Project [11]. Since then a number of studies have leveraged the crowd to *ask for help* [14, 35, 36]. Prior works have also built frameworks for web-based low-latency robot control [37]. Kehoe et al.[18] provides a treatise that touches on the aspects of cloud robotics: big data, cloud computing, collective robot learning and crowdsourcing. Teleoperation mechanisms vary from game interfaces [26] to free-space positioning interfaces [20]. A comparison of various control interfaces shows that general purpose hardware is deficient while special purpose hardware is more accurate but is not widely available [19, 26].

Virtual reality-based free-space controllers have recently been proposed both for data collection [23, 39] and policy learning [40, 43]. While these methods have shown the utility of data, they do not provide a seamlessly scalable data collection mechanism. Often the data is either collected locally or requires a powerful local client computer, to render the high definition sensor stream to a VR headset [39, 43]. The use of VR hardware and requirement of client-side compute resources has limited the deployment of these interfaces on crowdsourcing platforms.

Our system builds on RoboTurk [25], which uses a ubiquitous smartphone-based 6-DoF controller along with seamless cloud integration to ensure homogeneous quality of service regardless of client's compute resources. In contrast to local teleoperation methods that restrict data collection to a few users, crowdsourcing mechanisms such as RoboTurk can allow several interesting strategies to be demonstrated that vary across people, and across situations, leading to *diversity* of the data collected, as shown in Fig. 7.

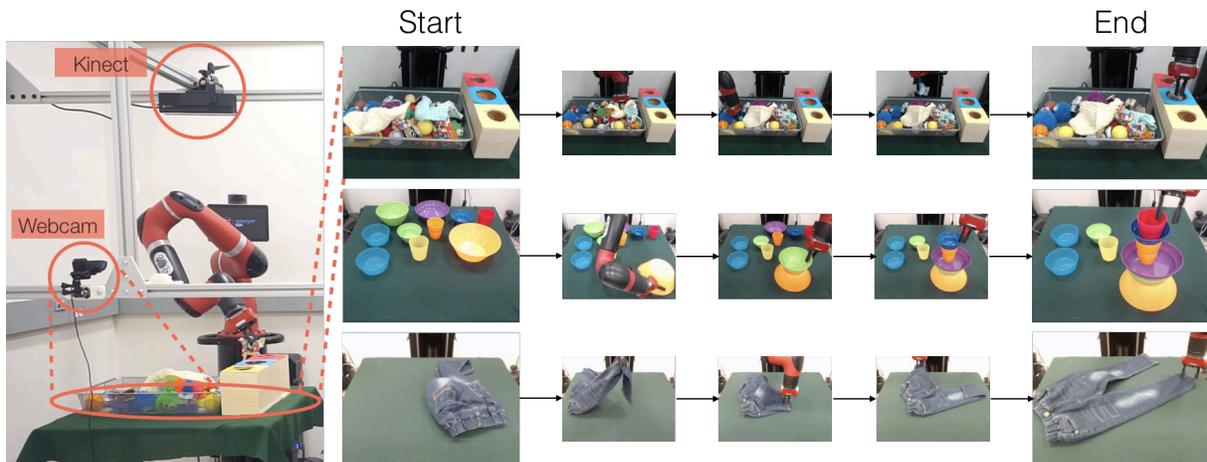


Fig. 3: **Data Collection Setup and Tasks:** Our data collection setup (left) consisted of three Sawyer robot arms, each of which had a front-facing webcam view and a top-down Kinect depth camera view. The front-facing view was streamed to each teleoperator. We collected data on three tasks (right) that require both fine-grained dexterity and high-level planning to solve. In the *Object Search* task (top) the objective is to find all instances of a certain target object category (plush animal, plastic water bottle, or paper napkin) and fit them into the corresponding bin. In the *Tower Creation* task (middle), the objective is to stack the various cups and bowls to create the tallest tower. In the *Laundry Layout* task (bottom) the objective is to layout an article of clothing on the table such that it lies flat without folds.

III. SYSTEM DESIGN

In order to collect our dataset, we leveraged the RoboTurk platform [25], which allows large groups of remote users to simultaneously collect task demonstrations by using their smartphones as motion controllers to control robots in simulated domains. We first review the original platform and then discuss the extensions we implemented to enable robust data collection on physical robot arms.

A. RoboTurk Overview

RoboTurk is a platform that allows users to seamlessly collect task demonstrations in simulation through low-latency teleoperation, regardless of their location or compute resources. Users connect to a website that streams video from the simulated environment, and use their smartphone as a motion controller to control the robot. The simulation itself runs in a remote server hosted in the cloud – this is to ensure homogeneous quality of service to every user regardless of available compute resources. In this way, RoboTurk facilitates crowdsourcing task demonstrations in simulated domains from large pools of remotely located annotators.

In order to support concurrent low-latency teleoperation servers that allow many users to use the platform simultaneously, the platform utilizes several core components. RoboTurk leverages Web Real-Time Communication (WebRTC) to establish low-latency communication links between a user’s phone, web browser, and a teleoperation server that is hosted in the cloud. We now outline the core components.

User Endpoint. Each user controls the robot arm using their smartphone as a motion controller and receives a real-time video stream of the robot workspace in their web browser. The phone transmits its changes in position and its current absolute orientation over WebRTC to the teleoperation server. We additionally provide users with the ability to enable and disable control of the robot so that users can re-position their

arm and phone for better ergonomics similar in nature to interfaces for surgical robots [9].

Coordination Server. This server creates and manages teleoperation sessions when users enter or leave the system.

Teleoperation Server. This is a process dedicated to a single user that is created by the centralized coordination server on a per-user basis. The server maintains its own simulator instance, sends frames to the user’s web browser, and handles the incoming control commands from the user.

B. Extending RoboTurk for use with real robots

While the initial RoboTurk platform works well for collection in simulated domains, collecting data on physical robots poses additional challenges.

Limited resources. While the number of simulation instances that can be run in parallel is bottlenecked only by compute resources, data collection on physical robots is bottlenecked by the number of available robots. Thus, to extend RoboTurk to the physical robot setting, we implemented a mutual exclusion principle to limit the number of users on the platform. The coordination server loads a centralized landing page for all users and routes them to a robot and corresponding task. It places a lock on the control of each physical robot arm so that only one user may operate the robot at a time.

Control latency. Controlling a robot in a simulator is markedly different from controlling a physical robot due to stochastic delays that can occur at a hardware level (e.g. commanding the motors) and at a network level, since commands are sent to the robot through a network connection. These are uncontrollable delays that are incurred in addition to those from the connection between the user and the teleoperation server. To account for these delays, we use a low-pass filter to reject high-frequency user input to ensure smooth and responsive teleoperation. By focusing on the low-frequency content in a user’s command stream, we are able to ensure that delays do not adversely affect a person’s

TABLE I: **Dataset Comparison.** We compare our dataset to similar robot datasets collected via human supervision in prior work. Items marked with * are estimates that were extrapolated using other reported information, and interfaces marked with † are not real-time.

Name	Interface	Task	Avg. Task Length (sec)	Number of Demos	Total Time (hours)
JIGSAWS[10]	daVinci	surgery	60*	103	1.66
Deep Imitation [43]	VR	pick, grasp, align	5*	1664	2.35
DAML[42]	Human demos	pick, place, push	5*	2941	4.08
MIME[34]	Kinesthetic	pick, place, push, pour	6*	8260	13.7*
PbD[8]	GUI†	pick, place	207*	465	25.8*
Roboturk-Real (Our)	iPhone AR	long horizon object manip	186	2144	111.25

capability to control the robot arm. Furthermore, instead of having teleoperation servers run on cloud infrastructure, we now spawn teleoperation servers on machines that are located in close physical proximity to the robots in order to minimize the latency of control commands sent to the robot arms.

Robot safety. We address the need for safety of the robot arms and workspaces by extending the RoboTurk smartphone app to ensure that the user is holding the phone correctly and exercising slow, deliberate motions by validating phone poses. Participants in our data collection study were given a structured 5 minute tutorial to familiarize themselves with operation of the physical robot arm, and also given a direct line of communication to people monitoring the robots to ensure quick responses to unsafe user control.

Data collection. Data collection in simulation is straightforward, since the state of the simulator can be saved at each timestep. From this minimal simulator state, all ground-truth observations of interest can be reconstructed. In contrast, data collection in the real world is much more unstructured. Multiple sensor streams emit data at different rates. To account for this, we leverage rosbag, a package built on top of the Robot Operating System (ROS), that allows for recording messages sent over concurrent publisher-subscriber communication channels called topics.

IV. DATA COLLECTION

A. Task Design

We designed three robotic manipulation tasks for data collection. These tasks were chosen with care - each task requires the use of low-level dexterity and high-level reasoning in order to solve it - both of which can be provided with a human in the loop. Furthermore, the solution space is multimodal - there are several admissible ways to solve the task for a given task instance. Consequently, there is inherent freedom in our tasks, encouraging our diverse set of users to experiment with different solution strategies.

Object Search. The goal of this task is to search for a set of target objects within a cluttered bin and fit them into a specific box. There are three target object categories: *plush animals*, *plastic water bottles*, and *paper napkins*. The workspace consists of a large cluttered bin containing a diverse assortment of objects and three target boxes, one per category of target object. At the start of each task instance, three target objects of each category are mixed in among the clutter of the box. A target category is randomly sampled and relayed to the operator, who must use the robot arm to find all three objects corresponding to the target category and place each item into its corresponding hole. We further place the constraint that

objects can be grasped and moved around within the bin but they cannot be placed outside the bin in any sort of staging area - this adds to the challenging nature of the task.

The Object Search task requires human-level reasoning to detect and search for the target items and dexterous manipulation to dig through the bin, push objects out of the way, pick up the target object successfully, and fit the target object into the corresponding hole, making it a good candidate for crowdsourcing. The objects also have interesting properties - the paper napkins appear in crumpled and unfolded configurations, and the crushed plastic water bottles are challenging to detect and grasp due to their translucence and arbitrary rigid shape. Furthermore, it is a practical problem with industrial applications [6].

Tower Creation. In this task, an assortment of cups and bowls are arranged on the table. The goal of the task is to create the tallest tower possible by stacking the cups and bowls on top of each other. This task requires physical reasoning over the properties of each type of cup and bowl and thinking about how to stack them on top of each other to maximize height without sacrificing the stability of the tower.

We diversify the initial task configurations by sampling a set of ten objects drawn from a total of 28 bowls in 7 varieties and 12 cups in 3 varieties. We also randomize the initial configuration of the objects. This encourages diversity in the demonstrations since users will not receive the same set of objects in the same configuration, enforcing each demonstration to be unique.

Laundry Layout. This task starts with a hand towel, a pair of jeans, or a t-shirt placed on the table. The goal is to use the robot arm to straighten the item so that it lies flat on the table with no folds. On every task reset we randomly place the item into a new configuration. This task was chosen for the visual and physical reasoning skills necessary to unfold and flatten the item. Solving it requires understanding the current item configuration and how it will respond to different types of contact.

B. Data Collection and Dataset Details

We collected our dataset using 54 different participants over the course of 1 week. Every user participated in a supervised hour of remote data collection, including a brief 5 minute tutorial at the beginning of the session. Afterwards, they were given the option to collect data without supervision for all subsequent collection. The users who participated in our data collection study collected the data from a variety of locations. All locations were remote - no data collection occurred in front of the actual robot arms.

Task statistics by previous trials of experience

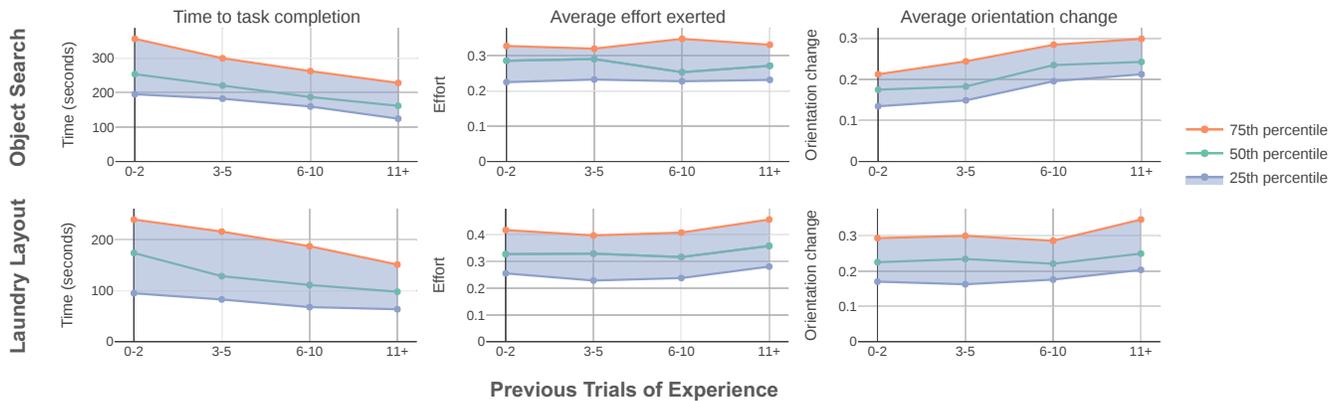


Fig. 4: **Characterizing user skill improvement over time.** Task completion times across users versus number of demonstrations of experience (left), average effort exerted versus experience (middle), and average change in orientation versus experience (right) on Object Search (top) and Laundry Layout (bottom). Together, these show that with experience, users learn to use more nuanced motion to complete tasks faster and more efficiently. Task completion time drops steadily as experience increases. However, average effort exerted (measured by the square L_2 norms of phone translations) is largely invariant across experience, while phone orientation change increases with experience, implying that users improve over time not by moving the phone faster, but rather by learning to enact more dexterous motions. Each graph displays quartiles to show that changes are consistent across the entire population. The Tower Creation task was excluded from this evaluation since most users insisted on using all five minutes to create their tower.

Fig. 3 shows our data collection setup. We collected data on three Sawyer robot arms - each of which had a front-facing webcam and a top-down Kinect depth camera mounted in the workspace of the robot arm. We collected RGB images from the front-facing RGB camera (which is also the teleoperator video stream view) at 30Hz, RGB and Depth images from the top-down Kinectv2 sensor also at 30Hz, and robot sensor readings such as joint and end effector information, at approximately 100Hz.

Table I compares our dataset against other robot datasets collected using humans. With over 111 hours of total robot manipulation data in our dataset, our dataset is 1-2 orders of magnitude larger than most other datasets. The number of task demonstrations in our dataset also compares favorably with the number of demonstrations in large datasets such as MIME [34], but the tasks that we collected data on are more difficult to complete, as they take on the order of minutes to complete successfully, as opposed to seconds (see Fig. 5).

V. EXPERIMENTS: SYSTEM ANALYSIS

A. Quantitative User Performance Analysis

The contributing users had a wide range of skill level, as seen in Fig. 5, but regardless of individual skill, we noticed strong trends of skill improvement with experience. We observed that users learned how to use the teleoperation control system very quickly, and once they became comfortable with it, they not only completed tasks faster, but more efficiently. Users demonstrated a higher degree of dexterity as they gained experience performing tasks.

Fig. 4 demonstrates quantitative results that corroborate what we observed. Task completion times steadily decreased as users gained experience doing the task. We also measured average user exertion, estimated by the square L_2 norms of phone translations during control, and found it largely

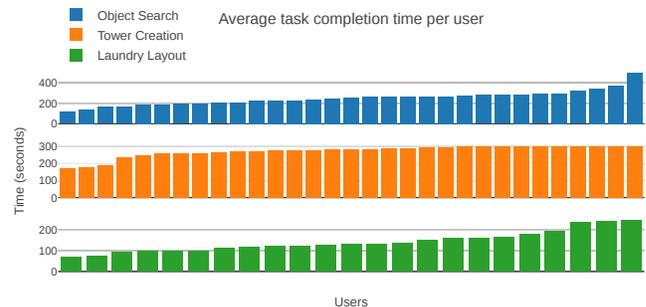


Fig. 5: **User skill comparison.** Average task completion time plotted for all users. The distribution indicates a wide range of skill exhibited across users on the Object Search and Laundry Layout tasks - more skillful users are able to complete these tasks in less time. By contrast, the uniform spread in times across the Tower Creation task indicates that all users were persistent in using all five minutes to create their tower.

invariant with respect to amount of experience, demonstrating that faster completion times were not due to users simply performing the same trajectory faster. By contrast, the average phone orientation change increased with experience, confirming that users learned to control the robot with more dexterous motion, enabling them to complete the task faster.

B. Qualitative User Feedback

Motivated by a previous study on robot teleoperation interfaces [19], we had each participant complete a NASA TLX form upon completion of the study [13]. This self-reported survey measured the participants' perception on mental demand, physical demand, temporal demand, performance, effort, and frustration on a 21-point scale. The total workload was computed as the sum of these averages, where higher scores represent a higher workload on users. From Table II, users found that the tower stacking task required the most workload across all the metrics.

TABLE II: NASA TLX Evaluation. Self-reported user evaluation of the system.

Measure	Object Search	Tower Stacking	Laundry Layout
<i>Mental demand</i>	12.2 ± 3.7	13.2 ± 4.4	9.3 ± 5.3
<i>Physical demand</i>	11.1 ± 4.0	11.8 ± 4.7	9.6 ± 4.6
<i>Temporal demand</i>	6.0 ± 4.6	12.2 ± 6.2	7.9 ± 4.9
<i>Performance</i>	6.3 ± 4.8	12.4 ± 5.1	7.1 ± 6.1
<i>Effort</i>	10.8 ± 4.2	14.2 ± 3.9	10.3 ± 5.1
<i>Frustration</i>	7.5 ± 5.0	13.1 ± 5.2	7.3 ± 5.4
Total Workload	53.9 ± 11.2	76.9 ± 12.2	51.5 ± 12.8

VI. EXPERIMENTS: DATA ANALYSIS

In this section, we investigate properties of our dataset that demonstrate the potential utility of the data for several applications such as multimodal density estimation, video prediction, reward function learning, policy learning and hierarchical task planning.

A. Evaluating Object Search Task Complexity

One of the key novelties of the dataset we present is the complexity of the reasoning necessary to plan a strategy for solving the tasks and the actual dexterity necessary to perform the finer details of manipulating the objects. We focus on the Object Search task since there is a simple qualitative measure to demonstrate the complexity of the task through the start and end configurations of the bin. Fig. 6 shows the start and end configurations of the bin over several successful task demonstrations. The fact that these configurations vary greatly is evidence that the operator needed to drastically reorient the contents of the bin to find the objects of interest.

B. Diversity of task demonstrations and solution approaches

One of the benefits of collecting data from a set of 54 users on tasks with an unstructured solution space is that every user has a unique approach to solve a given task instance. Fig. 7 presents three demonstrations on the Laundry Layout task that started with similar towel configurations and ended successfully with the towel flat on the table. The demonstration frames show that although the start and end configurations are similar, the approach used to flatten the towel are markedly different. For example, the user that provided the top demonstration chose to pick one side of the towel, place it down in a better configuration, then push the remaining flap of the towel open to complete the task. By contrast, the bottom demonstrator chose to grab a corner of the towel and manipulate it through lifting and dragging motions until the cloth was flat on the table.

Fig. 7 also presents three demonstrations on the Tower Creation task that started with the same set of objects, but resulted in three completely different towers. While both towers were comparably high, the users demonstrated different ways to stack the cups and bowls in order to build the tower. The dataset that we collected contains many such instances of multimodality and creative problem solving that stem from the diversity of the humans that generated the data.

C. Inferring a Reward Signal from Demonstrations

Consider the problem of learning a policy to imitate a specific video demonstration. Prior work [4, 33] has approached this problem by learning an embedding space



Fig. 6: **Object Search Task Complexity.** The large difference between start and end frames on the Object Search demonstrates the significant amount of effort required to solve the task.

over visual observations and then crafting a reward function to imitate a reference trajectory based on distances in the embedding space. This reward function can then be used with reinforcement learning to learn a policy that imitates the trajectory. Taking inspiration from this approach, we trained a modified version of Time Contrastive Networks (TCN) [33] on Laundry Layout demonstrations and investigate some interesting properties of the embedding space.

To address the large and diverse amount of data that was collected, we made two important modifications to the TCN algorithm. The original algorithm used a triplet loss to encourage neighboring video frames to be close in the embedding space; however, we found that applying the original TCN algorithm to our dataset resulted in embeddings with distances that were not meaningful for frames with larger time separation in a demonstration. Learning an embedding space that can tolerate frames with large temporal separation is critical for our dataset, since our tasks are multi-stage and our demonstrations are several minutes long.

In order to learn both high and low frequency temporal similarity, we split each demonstration into *chunks* of uniform size and use two separate triplet losses - an intra-chunk loss that pushes neighboring frames from within the same chunk of time together in the embedding space and an inter-chunk loss that encourages frames from nearby chunks of time to be close in the embedding space. We also added an auxiliary loss to encourage terminal demonstration frames to be close in the embedding space.

In Fig. 8, we consider the frame embeddings along a single Laundry Layout demonstration. We plot the negative L2 distance of the frame embeddings with respect to the embedding of a target frame near the end of the video, where the target frame depicts a successful task completion with the towel lying flat on the table. The figure demonstrates that distances in this embedding space with a suitable target frame yield a reasonable reward function that could be used to imitate task demonstrations purely from visual observations.

Furthermore, embedding distances capture task semantics to a certain degree and could even be used to measure task progress. For example, in frames 3 and 5, the towel is nearly flat on the table, and the embedding distance to frame 6 is correspondingly small. By contrast, in frames 2 and 4, the robot is holding the towel a significant distance away from the table, and the distance to frame 6 is correspondingly large.

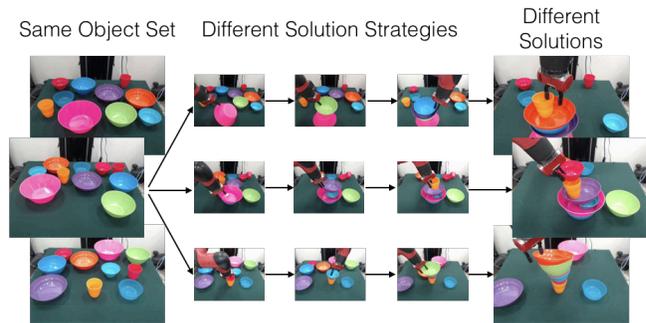
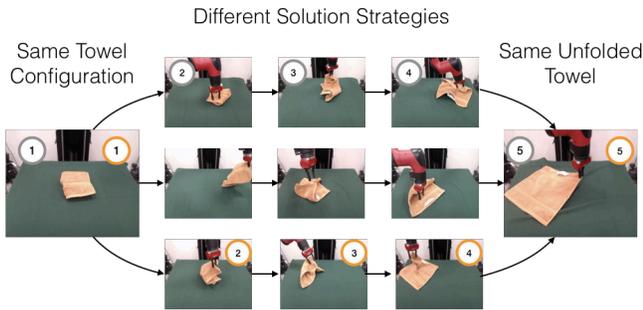


Fig. 7: **Diversity of Task Solutions.** Three different demonstrations provided on the Laundry Layout task (left) and the Tower Creation task (right) are shown above. The Laundry Layout demonstrations start and end with the same towel configuration across demonstrations but users exhibited different solution strategies to solve the task. The Tower Creation demonstrations start with the same set of objects in different configurations on the table but users chose to leverage the items in different ways, leading to three towers of roughly the same height, with different structural composition. This showcases the diversity of solution strategies that are present in our dataset.

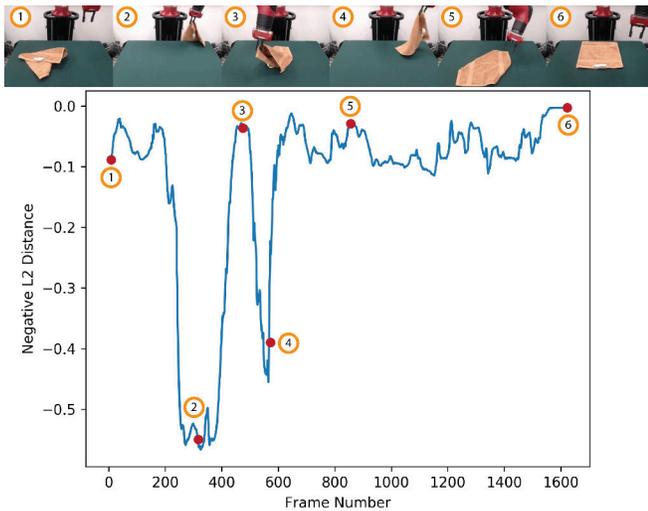


Fig. 8: **Leveraging learned embedding spaces as a similarity metric for imitation.** We trained a custom variant of TCN [33] to learn an embedding space over RGB images on the Laundry Layout task. We plot the negative L2 embedding distance between a target frame with a flat towel (the last frame) and all other frames in a demonstrations. This distance provides a meaningful reward function for imitation as well as a useful metric for task progress.

D. Behavioral Cloning

We also trained policies using Behavioral Cloning on the Laundry Layout task by learning a mapping from RGB images to robot joint positions. Our attempts to learn from the entire dataset were ultimately unsuccessful due to the diverse nature of the demonstrations, but we were able to achieve some success by restricting the training data to demonstration segments where the arm moves to a corner of the towel, and lifts the towel up. Addressing the diversity of the dataset for policy learning is left for future work.

VII. CONCLUSION

We introduced three challenging manipulation tasks: Object Search, Tower Creation, and Laundry Layout. Solving each of these tasks requires both higher-level reasoning to specify *what* to accomplish and dexterous manipulation that answers *how* to accomplish the necessary physical interactions. Each task instance also admits many diverse solutions, making

these tasks amenable to crowdsourcing.

We presented the largest known crowdsourced teleoperated robot manipulation dataset consisting of over 111 hours of data across 54 users. The dataset was collected in 1 week on 3 Sawyer robot arms using the RoboTurk platform. We made important extensions to the RoboTurk platform to enable data collection on physical robots, including accounting for additional delays in the remote teleoperation loop due to physical robot actuation, ensuring operational safety of the robots when being controlled by novices, and managing large-scale data collection across multiple sensor streams. We analyzed how the system allowed our participants to adapt quickly to the phone control interface in order to collect diverse, successful demonstrations on the three tasks. We also presented a set of qualitative and quantitative results that showcase the diversity and utility of the dataset. Our dataset can be useful for several applications such as multimodal density estimation, video prediction, reward function learning, policy learning and hierarchical task planning.

Future work will be focus on two directions - improving the platform and utilizing the dataset for policy learning. Platform improvements include: (1) alleviating the need for manual task resets by collecting data on reversible tasks (a forward task and a reset task) so that remote operators can reset the workspace and (2) developing a more structured scheduling system for operators that ensures fair waiting times for those in the queue. While our initial results suggest that the dataset can potentially be used for both one-shot imitation learning and direct imitation, further considerations and innovations will be necessary to handle the inherent diversity and multimodality of the solutions demonstrated.

Acknowledgement. Ajay Mandlekar is supported by the Department of Defense (DoD) through the NDSEG program. We acknowledge the support of ONR-MURI (1175361-6-TDFEK), Toyota Research Institute (S-2015-09-Garg), Weichai America Corp. (1208338-1-GWMZY), CloudMinds Technology Inc. (1203462-1-GWMXN), Nvidia (1200225-1-GWMVU), Toyota (1186781-31-UBLFA), Panasonic, and Google Cloud. We would like to thank the users that contributed to the dataset, as well as the Stanford People, AI & Robots (PAIR) group, and the anonymous reviewers for their constructive feedback.

REFERENCES

- [1] P. Abbeel, A. Coates, and A. Y. Ng, "Autonomous helicopter aerobatics through apprenticeship learning", *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1608–1639, 2010.
- [2] P. Abbeel and A. Y. Ng, "Inverse reinforcement learning", in *Encyclopedia of machine learning*, Springer, 2011, pp. 554–558.
- [3] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration", *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [4] Y. Aytar, T. Pfaff, D. Budden, T. Paine, Z. Wang, and N. de Freitas, "Playing hard exploration games by watching youtube", in *Advances in Neural Information Processing Systems*, 2018, pp. 2935–2945.
- [5] A. Boularias, J. Kober, and J. Peters, "Relative entropy inverse reinforcement learning", in *Proceedings of the 14th Int'l Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [6] M. Danielczuk, A. Kurenkov, A. Balakrishna, M. Matl, R. Martín-Martín, A. Garg, S. Savarese, and K. Goldberg, "Mechanical search: Learning manipulation strategies to extract a specific object from a bin of unknown objects", *The 2019 International Conference on Robotics and Automation (ICRA)*, 2019.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database", in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009.
- [8] M. Forbes, M. J.-Y. Chung, M. Cakmak, and R. P. Rao, "Robot programming by demonstration with crowdsourced action fixes", in *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.
- [9] L. Fritsche, F. Unverzag, J. Peters, and R. Calandra, "First-person teleoperation of a humanoid robot", in *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th Int'l Conference on*, 2015.
- [10] Y. Gao, S. S. Vedula, C. E. Reiley, N. Ahmidi, B. Varadarajan, H. C. Lin, L. Tao, L. Zappella, B. Béjar, D. D. Yuh, C. C. G. Chen, R. Vidal, S. Khudanpur, and G. D. Hager, "Jhu-isi gesture and skill assessment working set (jigsaws) : A surgical activity dataset for human motion modeling", 2014.
- [11] K. Goldberg, "Beyond the web: Excavating the real world via mosaic", in *Second International WWW Conference*, 1994.
- [12] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The columbia grasp database", in *Robotics and Automation, 2009. ICRA'09. IEEE International Conference on*, IEEE, 2009, pp. 1710–1716.
- [13] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research", in *Advances in psychology*, vol. 52, Elsevier, 1988, pp. 139–183.
- [14] P. F. Hokayem and M. W. Spong, "Bilateral teleoperation: An historical survey", *Automatica*, vol. 42, no. 12, pp. 2035–2057, 2006.
- [15] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks", in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [16] 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.
- [17] A. Kasper, Z. Xue, and R. Dillmann, "The kit object models database: An object model database for object recognition, localization and manipulation in service robotics", *The International Journal of Robotics Research*, vol. 31, no. 8, pp. 927–934, 2012.
- [18] 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.
- [19] D. Kent, C. Saldanha, and S. Chernova, "A comparison of remote robot teleoperation interfaces for general object manipulation", in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, ACM, 2017, pp. 371–379.
- [20] J. Kofman, X. Wu, T. J. Luu, and S. Verma, "Teleoperation of a robot manipulator using a vision-based human-robot interface", *IEEE transactions on industrial electronics*, 2005.
- [21] S. Krishnan, A. Garg, R. Liaw, B. Thananjeyan, L. Miller, F. T. Pokorny, and K. Goldberg, "Swirl: A sequential windowed inverse reinforcement learning algorithm for robot tasks with delayed rewards", *The International Journal of Robotics Research*, 2019.
- [22] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning hand-eye coordination for robotic grasping with large-scale data collection", in *ISER*, 2016, pp. 173–184.
- [23] J. I. Lipton, A. J. Fay, and D. Rus, "Baxter's homunculus: Virtual reality spaces for teleoperation in manufacturing", *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 179–186, 2018.
- [24] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojeda, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics", *arXiv preprint arXiv:1703.09312*, 2017.
- [25] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, S. Savarese, and L. Fei-Fei, "RoboTurk: A Crowdsourcing Platform for Robotic Skill Learning through Imitation", in *Conference on Robot Learning*, 2018.
- [26] R. Marín, P. J. Sanz, P. Nebot, and R. Wirz, "A multimodal interface to control a robot arm via the web: A case study on remote programming", *IEEE Transactions on Industrial Electronics*, 2005.
- [27] A. Y. Ng, S. J. Russell, et al., "Algorithms for inverse reinforcement learning.", in *Icml*, 2000, pp. 663–670.
- [28] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours", in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, IEEE, 2016.
- [29] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network", in *Advances in neural information processing systems*, 1989, pp. 305–313.
- [30] P. Rajpurkar, R. Jia, and P. Liang, "Know What You Don't Know: Unanswerable Questions for SQuAD", *arXiv preprint arXiv:1806.03822*, 2018.
- [31] S. Ross, N. Melik-Barkhudarov, K. S. Shankar, A. Wendel, D. Dey, J. A. Bagnell, and M. Hebert, "Learning monocular reactive uav control in cluttered natural environments", in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, IEEE, 2013, pp. 1765–1772.
- [32] J. Schulman, J. Ho, C. Lee, and P. Abbeel, "Learning from demonstrations through the use of non-rigid registration", in *Robotics Research*, Springer, 2016, pp. 339–354.
- [33] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, "Time-contrastive networks: Self-supervised learning from video", in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2018, pp. 1134–1141.
- [34] P. Sharma, L. Mohan, L. Pinto, and A. Gupta, "Multiple interactions made easy (MIME): large scale demonstrations data for imitation", *CoRR*, vol. abs/1810.07121, 2018. arXiv: 1810.07121.
- [35] A. Sorokin, D. Berenson, S. S. Srinivasa, and M. Hebert, "People helping robots helping people: Crowdsourcing for grasping novel objects", in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, IEEE, 2010, pp. 2117–2122.
- [36] J. Sung, S. H. Jin, and A. Saxena, "Robobarista: Object part based transfer of manipulation trajectories from crowd-sourcing in 3d pointclouds", in *Robotics Research*, Springer, 2018, pp. 701–720.
- [37] R. Toris, D. Kent, and S. Chernova, "The robot management system: A framework for conducting human-robot interaction studies through crowdsourcing", *Journal of Human-Robot Interaction*, 2014.
- [38] M. Večerík, T. Hester, J. Scholz, F. Wang, O. Pietquin, B. Piot, N. Heess, T. Rothörl, T. Lampe, and M. Riedmiller, "Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards", *arXiv preprint arXiv:1707.08817*, 2017.
- [39] D. Whitney, E. Rosen, E. Phillips, G. Konidaris, and S. Tellex, "Comparing robot grasping teleoperation across desktop and virtual reality with ros reality", *Int'l Symp. of Robotics Research*, 2017.
- [40] X. Yan, M. Khansari, Y. Bai, J. Hsu, A. Pathak, A. Gupta, J. Davidson, and H. Lee, "Learning grasping interaction with geometry-aware 3d representations", *arXiv preprint arXiv:1708.07303*, 2017.
- [41] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, "More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing", in *Int'l Conference on Intelligent Robots and Systems*, 2016.
- [42] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, and S. Levine, "One-shot imitation from observing humans via domain-adaptive meta-learning", *CoRR*, vol. abs/1802.01557, 2018. arXiv: 1802.01557.
- [43] T. Zhang, Z. McCarthy, O. Jow, D. Lee, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation", *arXiv preprint arXiv:1710.04615*, 2017.