

RoboSet

Vikash Kumar*, Rutav Shah^δ, Gaoyue Zhou^γ, Vincent Moens^φ,
Vittorio Caggiano^φ, Abhishek Gupta^ι, Aravind Rajeshwaran^φ

U.Washington^ι, UT-Austin^δ, CMU^γ, Meta-AI^φ

<https://sites.google.com/view/robohive>

2023

Motivation

For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

RoboHive offers extensive support for a diverse range of datasets known as *RoboSet*. These datasets are designed to facilitate pre-training and offline learning research for robotics. *RoboSet* was created to fill the gap that there is no large-scale robotics datasets in a uniform format that is suitable for robot learning.

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

RoboSet is created by researchers and students at the University of Washington, the University of Texas at Austin, Carnegie Mellon University, and Meta AI.

Who funded the creation of the dataset? If there is an associated grant, please provide the name of the

grantor and the grant name and number.

No specific funding agencies were involved.

Composition

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

Each instance of the dataset is a robot trajectory, capturing essential information at each time step. These trajectories comprise observations, actions, rewards, RGB visuals from multiple camera views, and other relevant environmental information.

RoboSet supports two types of trajectories: human teleoperation trajectories and expert trajectories. Human teleoperation trajectories are collected through human teleoperation using Puppet [1], while the expert trajectories are rollouts from a trained task-

*Correspondence to vikashplus@gmail.com

specific NPG policy for the target task. Please refer to Section *Collection Process* for data collection details.

How many instances are there in total (of each type, if appropriate)?

RoboSet spans across a diverse set of manipulation environments. For the specific number of instances available in each environment, please refer to [Table 1](#).

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set? If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).

RoboSet can be considered as a representative sample of robot manipulation trajectories in both simulated and real-world environments. It is important to note that this sample does not cover all possible manipulation tasks or robot morphologies, as the space of potential tasks is vast. Collecting trajectories for every possible task in all environments is currently impractical. However, we believe that the performance achieved on *RoboSet*'s baseline tasks can serve as a foundation for other manipulation tasks.

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images) or features? In either case, please provide a description.

Each instance, or trajectory, in *Ro-*

boSet contains the following information at each time step:

- Observations: information available to the robot, which contains the robot's joint state and any information relevant to completing the task in the environment.
- Actions: the robot's actions for each joint.
- Rewards: the reward the robot obtains.
- RGB visuals: RGB images from multiple camera views.
- Environment information: any relevant environment information, e.g. the state of particular objects in the environment, dense and sparse rewards, etc.

Is there a label or target associated with each instance? If so, please provide a description.

Every instance in *RoboSet* describes a robot trajectory. The actions of the trajectory can be seen as the label or target if the goal is to learn the robot's control policy.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

No. Since *RoboSet* is collected via the same platform, every trajectory contains the information mentioned above. For real-world data, logos over certain objects are blurred for anonymity and to adhere to the guidelines.

Are relationships between individual instances made explicit (e.g.,

Table 1: *RoboSet* data compositions across various domains and sources

Domain	# Trajs	# Tasks	World	Visuals	Source
Real kitchen	31000	40	Real	4 cam	Human TeleOp [1]
Franka kitchen	600	4	Sim	4 cam	Human TeleOp [2]
Adroit	25	4	Sim	3 cam	Human TeleOp [3]
Franka kitchen	75	5	Sim	4 cam	Expert Policy [4]
Adroit	75	4	Sim	3 cam	Expert Policy [4]
D’Kitty	75	3	Sim	4 cam	Expert Policy [4]

users’ movie ratings, social network links)? If so, please describe how these relationships are made explicit.

Yes. We group the dataset into different files indicating how the data was collected (human teleoperation or expert policy rollouts) and which environment the data was collected from.

Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.

Users have the flexibility to design their own data splits based on their specific needs when working with *RoboSet*. Additionally, users have the option to utilize all available trajectories for training, without the need to hold out a separate validation or testing set. In such cases, the trained policy can be directly evaluated within the simulator environment.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.

For simulated trajectories, rounding errors can potentially introduce noise to the data. However, we use float64 as the data type which offers higher precision and makes this noise negligible. For the real-world dataset, the noise

inherent in the hardware sensors is also represented in our dataset.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

The dataset is self-contained.

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals non-public communications)? If so, please provide a description.

Since some of the RoboHive environments are built upon existing environments such as Adroit and Franka Kitchen, it is natural to integrate datasets collected from the original environments. For instance, our human dataset of

the hand manipulation suite is adapted from the human trajectories collected from the DAPG project [5]. This enables us to reuse datasets from prior work effectively while supporting information that wasn't contained in the original dataset. For those trajectories, we comply with the original license and have also taken explicit consent from the authors to reuse those data.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

No. *RoboSet* consists solely of robot manipulation trajectories, which do not involve content that raises such concerns.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Collection Process

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

For expert datasets, we use a trained task-specific NPG policy for the target task and roll out 25 trajectories in the environment. We collect trajectories using three policies trained using three different seeds. The expert dataset may also contain failure trajectories.

Our human dataset is collected through human teleoperation using Puppet [1]. During the collection process, a human teleoperator uses an HTC Vive headset and controller to control the robot in an end effector space. We subsequently replay and parse the trajectories in each target environment to collect task-relevant information. The human trajectories in *RoboSet* are mostly successful.

For the trajectories we integrated from past projects, we replay the original trajectories into RoboHive's corresponding environments and record the trajectory with complete information that may or may not contain in the original trajectories, e.g. RGB observations.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)? How were these mechanisms or procedures validated?

The simulated trajectories in *RoboSet* is collected via the Mujoco [6] simulator. For collecting human trajectories, we use a virtual reality headset and controller to control the robot in an end effector space.

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

No, the dataset is not sampled from a larger set.

Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?

The dataset was collected by the authors of RoboHive. Two lab contractors responsible for the general upkeep of the robotics lab also helped with real-world data collection as part of their daily responsibilities.

Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.

RoboSet was collected intermittently over a span of two months between q3-2022 to q2-2023.

Were any ethical review processes conducted (e.g., by an institutional review board)? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

No, there was no need for ethical review as each instance is a robot manipulation trajectory.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

No.

Preprocessing/cleaning/labeling

Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)? If so, please provide a description. If not, you may skip the remainder of the questions in this section.

We parse *RoboSet* into HDF5 format, which is suitable for organizing large and complex hierarchical data.

Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the “raw” data.

No, because the raw data in pickle format contain the same information as the HDF5 files that we release.

Is the software used to preprocess/clean/label the instances available? If so, please provide a link or other access point.

Yes. The dataset processing script is available at RoboHive’s [codebase](#) on GitHub.

Uses

Has the dataset been used for any tasks already? If so, please provide a description.

We use the dataset in our baseline experiments, which compare the performance of various policy learning methods and visual representations.

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

<https://sites.google.com/view/robohive>

What (other) tasks could the dataset be used for?

The availability of both expert and human datasets in *RoboSet* makes it suitable for various learning approaches, encompassing not only imitation learning methods that rely on high-quality data, but also offline reinforcement learning

methods and any techniques that leverage multimodal or play data. Moreover, the wealth of visual data within *RoboSet* opens up avenues for visual pre-training and generalization.

Is there anything about the composition of the dataset or the way it was collected and pre-processed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

No.

Are there tasks for which the dataset should not be used? If so, please provide a description.

No.

Any other comments?

Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

Yes, the dataset is publicly available for anyone interested in utilizing it.

How will the dataset will be distributed (e.g., tarball on the website, API, GitHub) Does the dataset have a digital object identifier (DOI)?

The dataset is available as a compressed file for download on our website.

When will the dataset be distributed?

Our dataset is open-source and currently public. Given the size of the dataset, We are in the process of uploading the data. Majoring of the data has already been uploaded, we expect the full upload to finish by June'23

Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

The dataset will be released with license [CC-BY 4.0](#).

Have any third parties imposed IP-based or other restrictions on the data associated with the instances?

If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

No.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?

If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

No.

Any other comments?

Maintenance

Who will be supporting/hosting/maintaining the dataset?

The project has facilitated multiple projects over the years across multiple labs. Given the history of existing maintenance, we expect it to continue. The dataset will continue to be maintained by the authors of RoboHive.

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

Vikash Kumar, vikashplus@gmail.com

Is there an erratum? If so, please provide a link or other access point.

Please refer to our [release page](#).

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

Yes, we anticipate that *RoboSet* will continue to expand in terms of both quantity and diversity, ensuring its ongoing relevance and value to the research community. We expect to update the dataset as needed with new releases of RoboHive on GitHub.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

The dataset is not related to people.

Will older versions of the dataset continue to be supported/hosted/maintained? If so,

please describe how. If not, please describe how its obsolescence will be communicated to users.

Yes. Each version of the dataset will be tagged and released on GitHub.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description.

Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

To facilitate extensions we are providing extensive details, utilities, and tools to contribute to RoboSet.

Any other comments?

References

- [1] Vikash Kumar and Emanuel Todorov. Mujoco haptix: A virtual reality system for hand manipulation. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 657–663. IEEE, 2015.
- [2] Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long horizon tasks via imitation and reinforcement learning. *Conference on Robot Learning (CoRL)*, 2019.
- [3] Aravind Rajeswaran*, Vikash Kumar*, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [4] Sham M Kakade. A natural policy gradient. In *NIPS*, 2002.
- [5] Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. Learning Complex Dexterous Manipulation with Deep Reinforcement Learning and Demonstrations. In *Proceedings of Robotics: Science and Systems (RSS)*, 2018.
- [6] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, 2012.