

SAMPLE-HD: Simultaneous Action and Motion Planning Learning Environment

Michal Nazarczuk Tony Ng Krystian Mikolajczyk
Imperial College London

[michal.nazarczuk17, tony.ng14, k.mikolajczyk}@imperial.ac.uk

Abstract

Humans exhibit incredibly high levels of multi-modal understanding - combining visual cues with read, or heard knowledge comes easy to us and allows for very accurate interaction with the surrounding environment. Various simulation environments focus on providing data for tasks related to scene understanding, question answering, space exploration, visual navigation. In this work, we are providing a solution to encompass both, visual and behavioural aspects of simulation in a new environment for learning interactive reasoning in manipulation setup. *SAMPLE-HD* environment allows to generate various scenes composed of small household objects, to procedurally generate language instructions for manipulation, and to generate ground truth paths serving as training data.

1. Introduction

With the increasing interest in developing systems [4, 7, 11] capable of human like perception and reasoning capabilities, the amount of data required to develop data-driven approaches is constantly increasing and various synthetically generated datasets were introduced [1, 2, 10, 15]. Similarly, robotics community frequently uses simulated environments [3, 12] as means of training and testing algorithms due to the high cost of gathering data using the real hardware. However, synthetic environments that offer accurate physics simulation usually suffer from low visual quality [8, 13]. In this work, we are trying to bridge the gap between vision and robotics by proposing a new environment providing precise physics calculation along high quality visual output in order to facilitate training of reasoning algorithms in object manipulation setup.

Robotic manipulators are one of the main topics of interest in the robotics community. However, obtaining data to train neural networks with visual input involves either a time consuming acquisition process (restricted by real time operations) or training with the use of simulated environment. The latter often suffers from unrealistic visual quality making it less interesting from computer vision perspective.

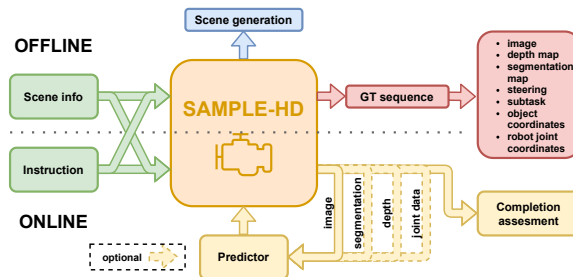


Figure 1. An overview of SAMPLE environment. Dotted line indicates offline and online capabilities. Different outward arrows show different steps of data generation or possible outputs.

SAMPLE-HD is a simulation environment build in Unity [14] that provides high quality visual data characterised by realistic physics behaviour. Thus, it can be used for tasks of scene understanding and controlling the robotic manipulator. We focus the simulation on a kitchen tabletop scenario, in which the robotic system is asked to perform object manipulation instructed by a given natural language query.

SAMPLE-HD provides a set of utilities to train visual systems to predict a steering signal for a robotic arm. An overview of the environment is presented in Figure 1. Firstly, it can generate various scenes, in which kitchen appliances (based on SHOP-VRB [9]) are randomly placed on the table with a randomised appearance. Additionally, we provide a way to procedurally generate a set of instructions asking the system to perform various interactions with the scene. The offline generated training data includes also a sequence of ground truth images and steering values required to fulfil the task. For the inference, or online training, we utilise ML-Agents [6] toolkit providing an API to facilitate communication between environment and the network.

Due to its modular nature, SAMPLE-HD allows for easy modifications. All the models used can be swapped for any imported 3D model which can conveniently be customised with respect to randomisation options. With the use of constrained-based templates, modifications to the instruction generator are also viable. Finally, ground truth manipulator paths are also based on simple modules allowing to stack them differently in order to obtain desired result.




Figure 2. Part of a GT sequence generated in SAMPLE-HD. Images are cropped, and sampled every 14th frame for better presentation.

2. Environment

Scenes in SAMPLE-HD are placed in virtual kitchen environment. All the objects and the manipulator are placed on the table. Robotic arm chosen for the simulation is Universal Robots UR5e equipped with a parallel gripper.

Steering of the robotic arm can be realised in different modes: in joint coordinates or cartesian space realised with inverse kinematics model. The control of the robotic arm is implemented asynchronously to physics calculation allowing to simulate vision-based decision making while preserving continuous movement. Self-collisions, collisions between robot and tabletop, and excessive forces are triggering safety mechanism stopping the arm movement.

Scene generation in SAMPLE-HD is fully randomised. Given the set of assets, each scene is generated via subsequent placement of objects in the scene. Every object is chosen randomly and its properties are also randomised. Via a menu attached to the object prototype, the user can choose which properties, among colour, material, size, of the asset can be changed. When placing the object, the possibility of intersections and occlusions with other objects are checked and incorrect positions are rejected. Finally, the scene is generated and rendered with the view from randomly positioned camera, segmentation map, depth map, ground truth objects data (in the manner of CLEVR annotations), and manipulator position.

Instructions asking the robotic arm to perform various actions related to manipulations are generated on the base of templates similar to CLEVR. We have noticed that when using depth first search on all possible adjectives combinations (like in CLEVR) leads to high execution times due to exponential increase with the number of adjectives. Therefore, we propose a rejection based generation of the instructions. Firstly, a random template is chosen assuring uniform distribution of either templates or resulting task types. For a chosen template and scene we generate a set of object level descriptors, depending on the required type - unique or non-unique. Descriptors are generated to provide a shortest unique (non-) description of the scene objects. In such way a more natural descriptions are created - it is more convenient to say *blue bowl* instead of *light, medium-sized, blue, hemispherical bowl*. Thereafter, we perform constraint based rejection test in order to filter out descriptions related to objects not fulfilling template criteria. Fi-

nally, filtered combinations are used to randomly choose set of descriptions and generate instructions. Additionally, we provide CLEVR-compatible programs for each instruction that can be used for symbolic approaches like [5, 16]. Finally, a descriptor of task along with pointers to target scene objects are included.

Motion paths generation is provided to obtain motion paths for generated instructions along with corresponding values of steering that may serve as the supervision for various data-driven approaches. Given the scene and the corresponding instruction, the task is split into major movements, namely: approaching the object’s neighbourhood, assuming grasping position, performing grasping, movement with object to destination, releasing the object.

Dataset utilising SAMPLE-HD is provided along the environment. We have chosen to use objects from SHOP-VRB. We provide 2000 scenes containing 3 to 5 objects. For each scene we attempt to generate instructions from provided templates resulting in 8360 instructions. Finally, a ground truth motion path was recorded at every rendered frame. The output consists of images rendered by a camera, segmentation and depth maps, robot position in joint coordinates, motion subtask and steering vector, and positions and orientations of every object in the scene (in such way the whole scene can be reconstructed if necessary). Additionally, SAMPLE-HD is equipped with a tool for assessing whether the task was fulfilled successfully. Example part of generated sequence is shown in Figure 2. As a part of dataset, we generated ground truth paths for all 8360 instructions that resulted in providing nearly 4.5 mln data points for training.




Figure 3. Example scene (left) from proposed dataset along with generated segmentation map (top right), depth map (bottom right), and sample instruction (bottom right).

References

- [1] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-Language Navigation: Interpreting visually-grounded navigation instructions in real environments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [2] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D Data in Indoor Environments. In *International Conference on 3D Vision (3DV)*, 2017. 1
- [3] Linxi Fan, Yuke Zhu, Jiren Zhu, Zihua Liu, Orien Zeng, Anchit Gupta, Joan Creus-Costa, Silvio Savarese, and Li Fei-Fei. SURREAL: Open-Source Reinforcement Learning Framework and Robot Manipulation Benchmark. In *Conference on Robot Learning (CoRL)*, 2018. 1
- [4] Haoshuo Huang, Vihan Jain, Harsh Mehta, Alexander Ku, Gabriel Magalhaes, Jason Baldridge, and Eugene Ie. Transferable Representation Learning in Vision-and-Language Navigation. In *International Conference on Computer Vision (ICCV)*, 2019. 1
- [5] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Inferring and Executing Programs for Visual Reasoning. In *International Conference on Computer Vision (ICCV)*, 2017. 2
- [6] Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. Unity: A General Platform for Intelligent Agents. *CoRR*, 2020. 1
- [7] Wonjae Kim, Bokyoung Son, and Ildoo Kim. ViLT: Vision-and-Language Transformer Without Convolution or Region Supervision. In *International Conference on Machine Learning (ICML)*, 2021. 1
- [8] A.T. Miller and P.K. Allen. GraspIt! *IEEE Robotics & Automation Magazine*, 11(4), 12 2004. 1
- [9] Michal Nazarczuk and Krystian Mikolajczyk. SHOP-VRB: A Visual Reasoning Benchmark for Object Perception. In *International Conference on Robotics and Automation (ICRA)*, 4 2020. 1
- [10] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, Dhruv Batra, and Facebook Ai Research. Habitat: A Platform for Embodied AI Research. In *International Conference on Computer Vision (ICCV)*, 2019. 1
- [11] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. ALFRED A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [12] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego De, Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy Lillicrap, and Martin Riedmiller. DeepMind Control Suite. *CoRR*, 2018. 1
- [13] Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 10 2012. 1
- [14] Unity-Technologies. Unity3D, 2020. 1
- [15] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building Generalizable Agents with a Realistic and Rich 3D Environment. In *Computer Vision - ECCV*, 2018. 1
- [16] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Kohli Pushmeet, and Joshua B Tenenbaum. Neural-Symbolic VQA: Disentangling Reasoning from Vision and Language Understanding. In *Advances in Neural Information Processing Systems*, 2018. 2