

iGibson 1.0: A Simulation Environment for Interactive Tasks in Large Realistic Scenes

Bokui Shen*, Fei Xia*, Chengshu Li*, Roberto Martín-Martín*, Linxi Fan, Guanzhi Wang, Claudia Pérez-D'Arpino, Shyamal Buch, Sanjana Srivastava, Lyne Tchapmi, Micael Tchapmi, Kent Vainio, Josiah Wong, Li Fei-Fei, Silvio Savarese

Abstract—We present *iGibson 1.0*, a novel simulation environment to develop robotic solutions for interactive tasks in large-scale realistic scenes. Our environment contains 15 fully interactive home-sized scenes with 108 rooms populated with rigid and articulated objects. The scenes are replicas of real-world homes, with distribution and the layout of objects aligned to those of the real world. *iGibson 1.0* integrates several key features to facilitate the study of interactive tasks: i) generation of high-quality virtual sensor signals (RGB, depth, segmentation, LiDAR, flow and so on), ii) domain randomization to change the materials of the objects (both visual and physical) and/or their shapes, iii) integrated sampling-based motion planners to generate collision-free trajectories for robot bases and arms, and iv) intuitive human-iGibson interface that enables efficient collection of human demonstrations. Through experiments, we show that the full interactivity of the scenes enables agents to learn useful visual representations that accelerate the training of downstream manipulation tasks. We also show that *iGibson* features enable the generalization of navigation agents, and that the human-iGibson interface and integrated motion planners facilitate efficient imitation learning of human demonstrated (mobile) manipulation behaviors. *iGibson 1.0* is open-source, equipped with comprehensive examples and documentation. For more information, visit our project website: <http://svl.stanford.edu/igibson/>.

I. INTRODUCTION

Simulation environments have proliferated over the last few years as a way to train robots and interactive agents in a rapid and safe manner. In these environments, agents learn to engage in physical interactions [1, 2], navigate based on sensor signals [3, 4, 5, 6, 7], or plan long-horizon tasks [8, 9, 10, 11]. In simulation, agents learn to perform interactions that actively change the input sensor signals and the state of the environment towards a desired configuration, capabilities at the core of what an embodied agent needs to achieve.

However, existing simulation environments that combine physics simulation and robotic tasks often cater to a narrow set of tasks and include only clean, small-scale scenes [12, 13, 14, 15, 16, 17]. The few simulation environments that include large scenes such as homes or offices either disable the possibility of interacting with the scene, focusing only on navigation (e.g. Habitat [18]), or use simplified modes of interaction (e.g. AI2Thor [19], VirtualHome [20]). These simulators do not support the development of end-to-end sensorimotor control loops for tasks that require rich, continuous interaction with the scene. Such tasks are difficult to accomplish in the aforementioned simulators, and the simplified modes of interaction lead to difficulties in transferring the learned policy onto real robots.

*Equal contribution. All authors are with the Stanford Vision & Learning Laboratory, Stanford University.

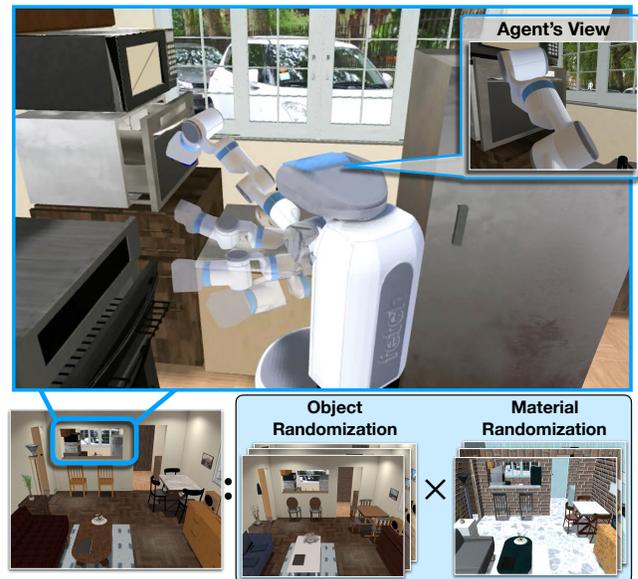


Fig. 1: Robot performs an interactive task in *iGibson 1.0*. It operates in the kitchen of one of *iGibson*'s fully interactive scenes, planning an interaction with the arm using an integrated sampling-based motion planner and receiving first-person view. *Bottom*: The same scene can be randomized with different materials and/or object models

We present *iGibson 1.0* (alternative called just *iGibson* in this manuscript), a novel simulation environment that enables the development of embodied agents for interactive tasks in large-scale realistic scenes (Fig. 1). Interactivity is achieved by leveraging a physics engine processing all elements in the scene, enabling manipulation of rigid and articulated objects as well as mobility. *iGibson 1.0* aims at unifying several aspects of robot simulation that are often available in different software tools, such as physics simulation for interaction with objects and robot control, high-quality simulated sensors, integration with reinforcement learning frameworks, and realistic indoor scenes that reflect the objects distribution of real homes. This integration allows fully physics based simulation of robot tasks (i.e. simulating the full complexity of the task) and allows developing task and motion planning, reinforcement learning or imitation learning solutions for those tasks with virtual sensor signals.

iGibson 1.0 contains 15 fully interactive and visually-realistic scenes with a total of 108 rooms. These scenes were generated by annotating 3D reconstructions of real-world scans (static scenes represented by a single mesh) and converting them into fully interactive scene models (scenes filled with articulated object models). In this process, we respect the original object-instance layout and object-

category distribution. The object models are extended from open-source datasets [21, 22, 15] enriched with annotations of material and dynamic properties. iGibson’s physics-based renderer leverages the extra information provided in the material annotation (maps of metallic, roughness and normals) to generate high-quality virtual images. To further facilitate the training of more robust visuomotor agents, iGibson 1.0 offers domain randomization procedures for materials (both visual appearances and dynamics properties) and object shapes while respecting the distribution of object placements and preserving interactability. iGibson 1.0 is also equipped with a graphical user interface that allows human users to easily interact with the scenes, enabling efficient collection of human demonstrations for imitation learning.

In summary, iGibson 1.0 provides the following novel features that facilitate developing and training robotic solutions:

- 1) Fifteen fully interactive scenes containing 108 rooms modelled after real world homes with articulated object models annotated with materials and dynamics properties. Additionally, we support importing CubiCasa5K [23] and 3D-Front [24] layouts, giving access to more than 12000 additional fully interactive scenes.
- 2) Realistic virtual sensor signals, including a physics-based renderer (PBR) for RGB images, rendering of normals, depth, point clouds, virtual LiDAR signals, and optical/scene flow. We further integrate domain randomization functionality (visual texture, dynamics properties and object instances) that facilitates generalization to unseen scenes.
- 3) Useful tooling for developing robotics solutions in simulation, such as a human-computer interface for humans to provide interactive demonstrations, and sampling-based motion planners for navigation and manipulation.

We demonstrate the benefits of these novel features in a comprehensive set of experiments in which visual agents are trained for navigation and interactive tasks. Our experiments show that iGibson 1.0 enables researchers to 1) train more robust and generalizable sensorimotor policies thanks to its realistic virtual sensor signals (including LiDAR) and domain randomization mechanisms, 2) collect human demonstrations and train imitation learning policies for manipulation and mobile manipulation tasks, and 3) learn intermediate visual representations linked to interactability of the scene that accelerate training of downstream manipulation tasks. iGibson 1.0 is open-source and academically developed, and available on our website <http://svl.stanford.edu/igibson/>.

II. RELATED WORK

The use of simulation in robotics has significantly increased in recent years and the research community has proposed a number of simulators for robotics and embodied AI. Here, we use the terms physics simulator and simulation environment as follows. A **physics simulator** is an engine capable of computing the physical effect of actions on an environment (e.g. motion of bodies when a force is applied, or flow of liquid particles when being poured) [27, 28, 29, 30, 31, 32]. On the other hand, a **simulation environment** is a framework that includes a physics simulator, a renderer of

virtual signals, and a set of assets (i.e. models of scenes, objects, robots) ready to be used to study and develop solutions for different tasks. Both components are crucial for advancing embodied AI and robotics. Here, we focus on the discussion of simulation environments.

Several simulation environments have been proposed recently to study manipulation with stationary arms [13, 12, 17, 33, 16]. Most of them are based on Bullet [27] or MuJoCo [28] for physics simulation, and render images with the default renderer or a Unity [30] plugin. Different from these simulation environments, iGibson focuses on large-scale (house-size) scenes and includes fifteen fully interactive scenes with 108 rooms – such as kitchens and bedrooms – where researchers can develop solutions for navigation, manipulation and mobile manipulation.

Closer to iGibson are simulation environments that include large-scale realistic scenes (e.g. homes or offices), which we summarize and compare in Table I. Gibson [25] (now Gibson v1) was the precursor of iGibson. It includes over 1400 3D-reconstructed floors of homes and offices with real-world object distribution and layout. Although Gibson incorporates PyBullet as its physics engine for simulating robot navigation, each scene is one single fully rigid object (static mesh). Thus, it does not allow agents to interact with the scenes other than collisions with the mesh, restricting its use to only navigation. A similar environment is Habitat [18]. Despite its high rendering speed, Habitat uses the non-interactive assets from Gibson v1 [25] and Matterport [34] and therefore only supports navigation tasks and simplified rearrangement of added objects. Recent work [35] introduced an extension to the Gibson v1 environment to support Interactive Navigation [35] where parts of the reconstructions corresponding to five object classes (chairs, tables, desks, sofas, and doors) in several Gibson static models were segmented and replaced with interactive versions. This enabled navigation agents to interact with the scene, and thus allowed for the first benchmark for Interactive Navigation. However, the simulators above are not interactive [25, 18] or partially interactive [35]. iGibson encapsulates the functionalities in all previous versions of Gibson with backwards compatibility for the static environments.

A variety of simulation environments have been proposed recently for scene-level interactive tasks, such as Sapien [15], AI2Thor [19], VirtualHome [20], and ThreeD-World (TDW [26]). These simulators adopt different ways of agent-world interactions. **Predefined Actions (PA)** consist in the set of actions that can be performed for each object type. When the agent is close enough to an object and the object is in the right state (precondition), the agent can select a predefined action, and the object is “transitioned” to the next state (postcondition). In Table I we refer to this technique as Rigid Bodies with Predefined Actions (RBPA). It is possible to combine Rigid Bodies with Predefined Actions (RBPA) and Rigid Body Physics (RBP), such as first using RBPA to grasp an object and then using RBP after releasing it. AI2Thor and VirtualHome use predefined actions (PA) as an abstraction of physical interactions and allow agents to

TABLE I: Comparison of Simulation Environments

	iGibson 1.0 (ours)	Gibson [25]	Habitat [18]	Sapien [15]	AI2Thor [19]	VirtualH [20]	TDW [26]
Provided Large Scenes Real-World / Designed	15 homes (108 rooms) / -	1400 / -	-	-	- / 120 rooms	- / 7	- / 25
Provided Objects Number / Materials	570 / Yes	- / -	- / -	2346 / No	609 / Yes	308 / No	200 / Yes
Agent/World Interaction Forces, Predefined Actions	F	-	-	F	F & PA	F & PA	F
Physics Engine	Bullet	Bullet	Bullet	PhysX	Unity	Unity	Unity & Flex
Type of Simulation Rigid Bodies Physics, Rigid Bodies with PA, Particles and Fluids	RBP	RBP	RBP	RBP	RBP&RBPA	RBPA	RBP&PF
Supported Task	Nav.&Manip.	Nav.	Nav.	Nav.&Manip.	Nav.&Manip.	Nav.&Manip.	Nav.&Manip.
Type of Rendering	PBR	IBR	PBR	PBR,RTX	PBR	PBR	PBR
Virtual Sensor Signals	RGB,D,N SS,FL,LiDAR	RGB,D N,SS	RGB,D,SS,S	RGB,D,SS	RGB,D,SS,S	RGB,D SS,FL	RGB,D,SS,S
Domain Randomization Scene,Object,Materials	S,O,M	-	-	-	S	S	S,O
Speed	++	+	+++	++(PBR)/-(RTX)	+	+	+
Human Interface	Mouse Keyboard	-	Mouse Keyboard	-	Mouse Keyboard	Natural Language	Virtual Reality
Integrated Motion Planner	Yes	No	No	No	No	No	No
Specialty	Phys. Int. in Large Scenes	Nav.	Fast, Nav.	Articulation, Ray Tracing	Object States, Task Planning	Object States, Task Planning	Audio, Fluids

Type of rendering: PBR:Physics-Based Rendering, IBR:Image-Based Rendering, RTX:Ray Tracing

Virtual sensor signals: RGB: Color Images, D:Depth, N:Normals, SS:Semantic Segmentation, LiDAR:Lidar, FL:Flow (optical and/or scene), S: Sounds

modify the object states instantaneously (e.g. open a closed cabinet), suitable to study high-level task planning. While the availability of PA helps focusing on symbolic reasoning, the full robotics problems require RBP to simulate the tasks in all its complexity. PA limits access to all granularities of the task, which impedes robot learning and sim2real transfer. With the purpose of simulating full robotics tasks, iGibson uses RBP, simulating the physics behavior of all objects continuously, with embodiment of real robots. This is crucial if the learned policy is to be deployed in the real world. Similar to iGibson, Sapiens also uses RBP without PA, but with smaller scenes, focusing on interaction with articulated objects. In iGibson, we enrich the articulated objects models from the PartNet-Mobility dataset introduced by Sapien with materials and dynamics properties. TDW is also capable of continuous RBP, with simplifications that facilitate grasping using robot avatars and not real robot platforms.

While other simulation environments focus on particular aspects of embodied agent simulation (Table. I), iGibson **uniquely** unifies a set of important tools for robotics that together enable robot learning in large scenes: support of LiDAR and PBR rendering, speed that enable reinforcement learning, robot integration (URDF, controllers, motion planners), and continuous physics simulation of agents and objects. This integration of features enables tasks such as mobile manipulation, illustrated in Fig. 2, including physics-based robot simulation across the entire task.

III. IGIBSON SIMULATION ENVIRONMENT

In this section, we discuss the main structure, properties and features of iGibson that support training of robust sensor-guided policies for navigation and manipulation.

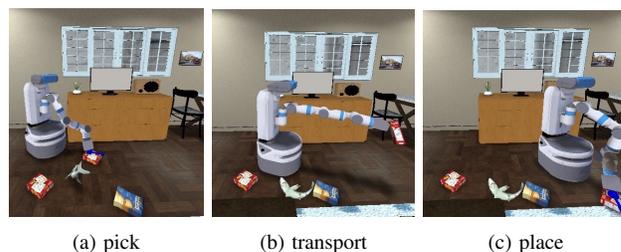


Fig. 2: Simulated Fetch Robot performing a mobile manipulation task in a cluttered environment.

A. Simulation Characteristics and API

At the highest level, iGibson follows the OpenAI Gym [36] convention. The environment receives an action and returns a new observation, reward and additional meta-information (e.g. if the episode has ended). Environments are specified with config files that determine scenes, tasks, robot embodiments, sensors, etc. Given a config file, iGibson creates an `Environment` that contains a `Task` and a `Simulator`. The `Simulator` contains a `Scene`, with a list of interactive `Objects` and one or more `Robot` instances. It also contains a `Renderer` that generates virtual visual signals from any point of view, such as a camera mounted on a robot or an external third person view. The `Task` defines the reward, initial and final conditions for the scene and the agents. While modular and easy to extend, most users may only need to interface with `Environment` after instantiating it with the appropriate config files.

iGibson comes with multiple easy-to-use configs, demos and Docker [37] files. It has been extensively adopted to train visuo-motor policies that successfully transfer to the real world [38, 39, 40, 41], and was the platform for iGibson Sim2Real Challenge at CVPR20 [42] and iGibson Challenge at CVPR21 [43]. The provided virtual LiDAR sensor has been used for robotics research in planning

and reinforcement learning for social navigation [44] and mobile manipulation [10]. iGibson is easily parallelizable and supports off-screen rendering on clusters.

B. Fully Interactive Assets

iGibson provides fifteen high quality fully interactive scenes with 108 rooms (see Fig. 3), populated with interactable objects. The scenes are interactive versions of fifteen 3D reconstructed scenes included in the Gibson v1 dataset. To preserve the real-world layout and distribution of objects, we follow a semi-automatic annotation procedure. This process is radically different from the annotation we performed for the Interactive Gibson Benchmark [35]. Instead of segmenting the original scene and replacing part of the meshes with interactive object models, we create fully interactive counterparts of the 3D reconstruction from scratch. This eliminates the need to fix artifacts in the original mesh due to reconstruction noise or segmentation error, and allows us to improve the overall quality of the scenes.

The scene generation process is composed of two annotation phases. First, the layout of the scene is annotated with floors, walls, doors and window openings. Then, all objects are annotated with 3D bounding boxes and class labels. We annotate bounding boxes for 57 different object classes, including all furniture types (doors, chairs, tables, cabinets, TVs, shelves, stoves, sinks, etc) and some small objects (plants, laptops, speakers, etc); see project website for the complete list. Annotating class-labeled bounding boxes allows us to scale and use different models of the same object class, while maintaining the real-world distribution of objects in the scene. In this way, we are able generate realistic randomized versions of the scenes (see Sec. III-D). To achieve the highest quality, for each class-labeled bounding box, we select a best fitting object model. The scene is also annotated with lights, with which we generate light probes for physics-based rendering (see Sec. III-C). We also bake in a realistic ray-traced ambient light and other light effects in the walls, floors and ceilings.

The object models are curated from open-source datasets: ShapeNet [21], PartNet Mobility [15, 45], and SketchFab. To preserve visual realism of the original reconstruction, we improve the object visual quality by annotating different parts of the models with photorealistic materials, which are then used by iGibson’s physics-based renderer. We utilize materials from CC0Texture, including wood, marble, metal, etc. To achieve a high degree of physics realism, we curate a mapping from visual materials to friction coefficients. We additionally compute the collision mesh, center of mass and inertia frame for each link of all objects. To assign realistic mass and density for different objects, we take the the median values of the top 20 search results from Amazon.

Additionally, we provide compatibility with CubiCasa5K [23] and 3D-Front [24] repositories of home scenes. We use their scene layouts and populate them with our annotated object models, leading to additional more than 12000 interactive home scenes. These scenes contain fewer

objects than the fifteen iGibson scenes, but provide a very large number of additional models to train tasks.

The fully interactive scenes we include in iGibson enable learning of interactive tasks in large realistic home scenes; in Sec. IV-C we show that the scenes can be used to learn a useful visual representation that accelerates the learning of downstream manipulation tasks.

C. Virtual Sensors

A crucial component of iGibson is the generation of high quality virtual sensor signals, i.e. images and point clouds, for the simulated robots. In the following, we summarize the most relevant of these signal generators (Fig. 4).

Physics Based Rendering: In iGibson, we include an open-source physics-based renderer, which implements an approximation of BRDF models [46] with spatially varying material maps including roughness, metallic and tangent-space surface normals, extending [47].

LiDAR Sensing: Many real-world robots are equipped with LiDAR sensors for obstacle detection. In iGibson, we support virtual LiDAR signals, with both 1 beam (e.g. Hokuyo) and 16 beams (e.g. Velodyne VLP-16). We include a simple drop-out sensor noise model to emulate the common failure case in real sensors in which some of the laser pulses do not return. Additionally, we provide the functionality to turn the 1D LiDAR scans into local occupancy maps, which are bird’s-eye view images with three types of pixels indicating free, occupied, or unknown space.

Additional Visual Channels: In addition to RGB and LiDAR, we support a wide range of visual modalities, such as depth maps, optical/scene flow and normals, segmentation of semantic class, instance, material and movable parts. These modalities can support research topics such as: depth/segmentation/normal/affordance prediction [48, 49, 50], action-conditioned flow prediction [51], multi-modal pose estimation [52, 53, 54], and visuomotor policy training assuming perfect vision systems [35, 55].

D. Domain Randomization

It is standard practice for robot learning to partially randomize the environment’s parameters in order to make the policy more robust [56, 57, 58, 59]. With the model being trained in a wide distribution of environments, it will be more likely to generalize to unknown evaluation environments. The evaluation environment may be the real world if we aim to train in simulation and transfer the policy to a real robot. In iGibson, we include domain randomization that leads to an endless variation of visual appearance, dynamics properties and object instances with the same scene layout.

First, we provide **object randomization**. The original 3D reconstructions are annotated with class-labeled object bounding boxes. These labels can be used to instantiate any object model of the corresponding class into the given bounding box (e.g. a bounding box labeled as “table” can be filled with any table model). This randomization maintains the semantic layout of the scenes (i.e. the object categories remain at the same 3D locations) while enabling near-infinite



Fig. 3: Fifteen interactive iGibson 1.0 scenes modeled after real-world reconstructions, preserving layout, distribution and size of objects.

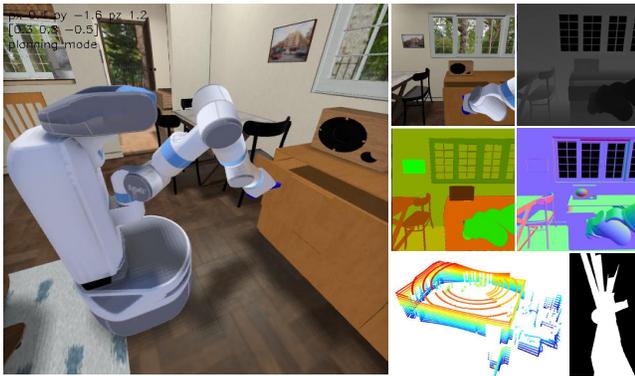


Fig. 4: Robot interacting in iGibson 1.0 (large picture: 3rd person view) and **virtual sensor signals** generated. Policies and solutions can make use of the following channels: (from top to bottom, left to right) RGB, depth, semantic/instance segmentation, normals, 16D LiDAR (point cloud), 1D LiDAR (also as occupancy map). Not depicted: optical/scene flow, joint encoders for robot’s and objects’ joints, poses, wrenches, contact points, and map localization.

combinations of object instances. It provides strong variation in depth maps and LiDAR signals that helps robustify policies based on these observations (see Sec. IV-A).

Second, we provide **material randomization**. In addition to high-quality material annotation for object and scene models, we provide a mechanism to randomize the specific material model associated with each object part (e.g. associating a different type of wood or metal). The effect is a stark color randomization that still represents plausible material combinations. This randomization generates strong variations in the RGB images and helps robustify policies based on this observation (see Sec. IV-A). Moreover, the dynamics properties of all object links can be randomized based on a curated mapping from visual materials to dynamics properties.

E. Motion Planning

Motion planners provide collision-free trajectories to move a robot from an initial to a final configuration [60]. They can be used to generate collision-free navigation paths for robot bases and collision-free motion paths for robot arms. In iGibson, we include implementations of the most popular sampling-based motion planners: rapidly growing random trees (RRT [61]) and its bidirectional variant (BiRRT [62]),

and lazy probabilistic road-maps (lazyPRM [63]), adapted from [64]. Sampling-based motion planners can have rather suboptimal and intricate paths. To alleviate this, we include acceleration-bounded shortcuts [65] for smoother paths.

F. Human-iGibson Interface

We provide a human-iGibson interface that enables users to navigate and interact in iGibson scenes using mouse and key commands on a viewer window. The user can navigate and interact with (pull, push, pick and place) objects. While a virtual reality or a 3D mouse interface may provide a more intuitive experience, most users do not have the necessary hardware. This interface offers a natural and simple way to demonstrations for imitation learning, evaluate the difficulty or feasibility of a task, or change the scene into a better initial state, for example. This interface is also integrated with the motion planner to command the robot to desired base and/or arm configurations. We verify this interface facilitates efficient development of interactive robotic solutions in Sec. IV-B .

IV. EXPERIMENTS

The goal of our experiments is to study how iGibson’s features help to develop AI agents. Specifically, we examine:

- (Sec.IV-A) does iGibson’s **domain randomization** and realistic virtual sensor signals (including **LiDAR**) allow navigation agents to generalize to unseen scenes (including the real world)?
- (Sec.IV-B) can the **human-iGibson interface** be used to efficiently train imitation learning agents for manipulation and **mobile manipulation tasks**?
- (Sec.IV-C) does the **full interactivity** in the scenes allow agents to learn visual representations that accelerate learning of downstream manipulation tasks?

A. Domain Randomization and Realistic Virtual Sensor Signals for Robot Navigation

In the first three experiments, we evaluate the generalization benefits brought by our realistic virtual sensor signals (including LiDAR) and domain randomization. First, we compare the generalization capabilities of vision-based reinforcement learning policies trained with and without domain randomization. Concretely, we evaluate the performance of



Fig. 5: Robot navigating the real-world counterpart of the iGibson 1.0 scene `Rs_int`. The robot executes a policy trained in simulation with virtual LiDAR signals, without domain adaptation. The quality and realism of iGibson facilitates zero-shot policy transfer.

policies trained in iGibson for PointGoal tasks [66] in held-out scenes with held-out visual textures. The observations for the policy include the depth maps, the robot’s linear and angular velocities, the goal location in the robot’s reference frame, and the next 10 waypoints in the shortest path between the robot’s current location and the goal location, separated by 0.2m. The waypoints are computed only based on the room layout, not the objects, so the robot mainly relies on depth maps for obstacle avoidance.

Second, we evaluate the performance of robot policies trained in iGibson to navigate to a target object (a lamp) using virtual RGB images also in held-out scenes with held-out visual textures. The task goal is to get at least 5% of the image occupied by the pixels of the target object. The observation for the policy only includes RGB images. In the above two experiments, we train in eleven scenes and evaluate in four held-out scenes with held-out visual textures.

Third, we evaluate the performance of the policies trained in iGibson for a PointGoal tasks [66] using **virtual LiDAR signals** with no vision inputs, and examine how well those policies transfer to the real world without adaptation, a hard test for generalization. The observations for the policy include a 1D LiDAR scan with 512 laser rays, the robot’s linear and angular velocities, and the goal location in the robot’s reference frame. To focus on sim2real transferability, we train in our scene `Rs_int`, for which we have access to the real-world counterpart (see Fig. 5).

Results: In the first two experiments, we observe better generalization capabilities in policies using iGibson’s domain randomization. For PointGoal navigation based on depth images, the performance goes from 0.27 to 0.40 SPL [66] and from 31.25% to 44.75% success rate when using randomization, indicating that the larger variety of shapes observed in the training process generates more robust depth-based policies. For object navigation based on RGB images, the performance goes from 49.75% to 57.5% success rate, indicating that material randomization helps in obtaining RGB-based policies that are more generalizable to unseen scenes and textures. Finally, for PointGoal navigation based on LiDAR signals, the policy achieves 33% success rate in `Rs_int` in iGibson, and 24% success rate in the real-world apartment. With only a 9% drop in performance and the failures mostly occurring in the same episodes (same pairs of the initial and goal locations in iGibson and real

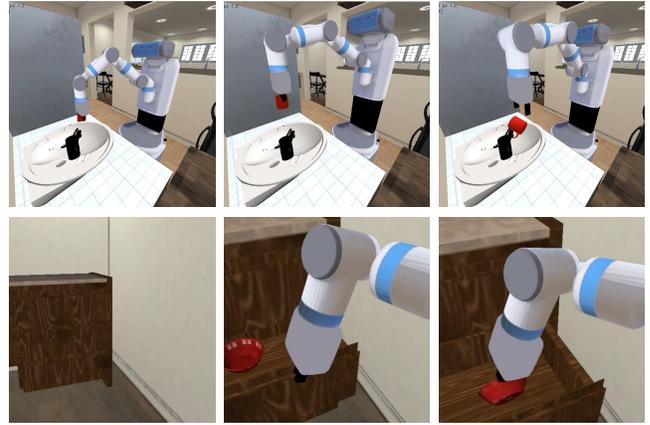


Fig. 6: Imitation learning from human demonstration. *Top:* third-person view of a robot performing a pick-and-place task. The policy trained using demonstrations collected with our human-iGibson interface achieves 98% success rate. *Bottom:* first-person view of a robot performing a mobile manipulation task. The policy trained using teleoperated demonstrations achieved 70% success rate.

world), this experiment indicates that the LiDAR signals generated in iGibson are realistic enough to facilitate zero-shot policy transfer. In summary, as shown in Table. I, iGibson provides unique support to train with realistic virtual sensor signals (e.g. LiDAR) and domain randomization, which leads to more robust robot navigation policies that successfully transfer to novel scenes.

B. Imitation Learning of Human Demonstrated (Mobile) Manipulation

In the second set of experiments, we evaluate iGibson as platform to train robots to perform manipulation and mobile manipulation tasks with Imitation Learning. First, we test the usability of the human-iGibson interface to efficiently collect demonstrations of **manipulation-only** tasks. We collect 50 demonstrations of pick-and-place operations with 20 mug models: pick a mug and place in the sink (Fig. 6), and store pairs of state (object position) and action (desired end-effector translation). We use the demonstrations to train a behavioral cloning policy that maps states to actions at 20 Hz. The action space consists of two parts: a 3-dimensional continuous space for desired end-effector translation, and a 1-dimensional discrete space for gripper opening. The desired end-effector translation is computed using inverse kinematics (IK) and executed with a joint position controller. The evaluation is conducted using a simulated mobile manipulator (Fetch robot), and generalization is tested with 5 unseen mugs.

Second, we collect demonstrations for imitation learning through continuous teleoperation using a system similar to Roboturk [67] for **mobile-manipulation** tasks. We collect 200 demonstrations with a simulated Fetch robot for search-and-pick operations: the robot must navigate and interact with a cabinet to open drawers, find a bowl and pick it (Fig. 6 (bottom)). The bowl and robot initial poses are randomized between episodes. Using these demonstrations, we train a behavioral cloning policy mapping observations to actions at 20 Hz. The observation space includes the robot

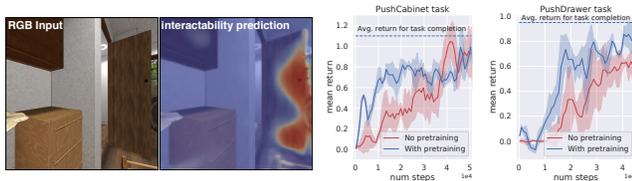


Fig. 7: *Left*: Example result of interaction pretraining. The agent receives RGB input and predicts if the pixels are pushable (red: higher probability; blue: lower) (Sec. IV-C). The model learns to associate edge of doors as the most pushable points. *Right*: Training curves for two interactive tasks (PushDrawer, PushCabinet) with and without interaction pretraining.

proprioceptive information (joint positions and velocities) and RGB-D images from virtual cameras on robot’s head and wrist. The action space consists of four parts: a 1-dimensional discrete value indicating whether to extend the arm, 2-dimensional continuous values representing the robot base’s linear and angular velocities, 6-dimensional continuous values representing the desired pose change of the end-effector, and 1-dimensional discrete value indicating whether to open the gripper.

Results: In our first experiment, training manipulation-only policies with imitation learning, we observe 98% success rate over 100 evaluation episodes. This experiment showcases that the human-iGibson interface enables easy collection of effective demonstrations for imitation learning, and the integrated motion planner is helpful for policy training. These two integrated features, as shown in Table. I, are novel combinations offered by iGibson. In the second experiment, training mobile-manipulation policies with imitation learning, we observe 70% success rate over 20 evaluation episodes. This experiment indicates that we can leverage iGibson to train imitation learning algorithms for mobile manipulation tasks.

C. Pretraining in Fully Interactive Scenes

In the third and final set of experiments, we evaluate the potential of using iGibson’s fully interactive scenes to learn an intermediate visual representation that encodes the expected outcome of interactions with different objects. Such an intermediate visual representation may be used to accelerate robot learning of manipulation tasks, since they typically require the agent to associate visual observations with promising areas of interaction to change the state of the scene towards a manipulation goal.

To learn such representations, we set up a virtual agent that interacts with random points in the scenes and learns to predict the outcome of these interactions. The interaction is parameterized as a coordinate in the virtual agent’s image observation space. We emulate a pushing interaction by displacing the corresponding 3D location of the selected pixel by 30 cm in the opposite direction of the surface normal, applying a maximum force of 60 N (a common payload of commercial robots). A motion of the point for more than 10 cm is considered a success. We sample 10 random pushes at each location, 4,000 locations in each scene. We use the images annotated with interaction successes/failures to train

a U-Net [68]-based visual encoder that predicts heatmaps of expected interaction success from RGB input.

For the second phase, we train two policy networks for two manipulation task respectively (PushDrawer, PushCabinet). The goal is to close the drawers or the cabinets. The policy outputs points to interact (push) that are given to one of our integrated motion planners to generate an arm motion [10]. We use DQN [69] as policy learning algorithm. The predicted interaction heatmaps are used to gate the Q-value maps predicted by the network.

Results: Fig. 7 (left) depicts the result of the pretrained visual model. We observe that the heatmap has stronger activation at the edge of the door than in the area closer to the hinge, and weak activation on closed cabinets. The model learns to identify the best areas to push to cause motion (further visualizations on project website). For both downstream tasks, we observe that using the pre-trained representation significantly accelerates training (Fig. 7 (right)). This suggests that the full interactability of iGibson can help agents learn useful visual representation for downstream mobile manipulation tasks. Having a subset of objects not physically interactable will lead to false negatives during pretraining, and prevents successful representation learning. As shown in Table. I and discussed in Sec. II, fully interactive scenes with continuous robot actions is a specialty of iGibson.

V. CONCLUSION

We presented iGibson, a novel simulation environment for developing interactive robotic agents in large-scale realistic scenes. iGibson includes 15 fully interactive scenes with 108 rooms, and novel capabilities to generate high-quality virtual sensor signals, domain randomization, integration with motion planners, and human-iGibson interface. Through experiments, we showcased that iGibson helps to develop robust policies for navigation and manipulation. We hope that iGibson can aid researchers in solving complex robotics problems in large-scale realistic scenes.

VI. ACKNOWLEDGEMENT

We thank NVIDIA, Google, ONR MURI (N00014-14-1-0671), ONR (1165419-10-TDAUZ), Panasonic (1192707-1-GWMSX), Qualcomm and Samsung for their support.

REFERENCES

- [1] T. P. Lillicrap *et al.*, “Continuous control with deep reinforcement learning,” *ICLR*, 2016.
- [2] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *JMLR*, 2016.
- [3] P. Mirowski *et al.*, “Learning to navigate in complex environments,” *ICLR*, 2017.
- [4] E. Parisotto and R. Salakhutdinov, “Neural map: Structured memory for deep reinforcement learning,” *ICLR*, 2018.
- [5] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, “Target-driven visual navigation in indoor scenes using deep reinforcement learning,” in *IEEE ICRA*, 2017.
- [6] W. B. Shen *et al.*, “Situational fusion of visual representation for visual navigation,” in *IEEE ICCV*, 2019.
- [7] K. Chen, J. P. de Vicente, G. Sepulveda, F. Xia, A. Soto, M. Vázquez, and S. Savarese, “A behavioral approach to visual navigation with graph localization networks,” *RSS*, 2019.
- [8] C. R. Garrett *et al.*, “Ffrob: Leveraging symbolic planning for efficient task and motion planning,” *IJRR*, 2018.

- [9] D. Xu *et al.*, “Neural task programming: Learning to generalize across hierarchical tasks,” in *IEEE ICRA*, 2018.
- [10] F. Xia, C. Li, R. Martín-Martín, O. Litany, A. Toshev, and S. Savarese, “Relmogen: Leveraging motion generation in reinforcement learning for mobile manipulation,” *IEEE ICRA*, 2021.
- [11] E. Li, R. Martín-Martín, F. Xia, and S. Savarese, “Hrl4in: Hierarchical reinforcement learning for interactive navigation with mobile manipulators,” in *CoRL*, 2019.
- [12] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, “Rlbench: The robot learning benchmark & learning environment,” *IEEE RA-L*, 2020.
- [13] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine, “Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning,” in *CoRL*, 2020.
- [14] Y. Urakami, A. Hodgkinson, C. Carlin, R. Leu, L. Rigazio, and P. Abbeel, “Doorgym: A scalable door opening environment and baseline agent,” *arxiv:1908.01887*, 2019.
- [15] F. Xiang *et al.*, “Sapien: A simulated part-based interactive environment,” in *IEEE CVPR*, 2020.
- [16] Y. Lee, E. S. Hu, Z. Yang, A. Yin, and J. J. Lim, “IKEA furniture assembly environment for long-horizon complex manipulation tasks,” *IEEE ICRA*, 2021.
- [17] Y. Zhu, J. Wong, A. Mandlekar, and R. Martín-Martín, “robosuite: A modular simulation framework and benchmark for robot learning,” *arxiv:2009.12293*, 2020.
- [18] Manolis Savva*, Abhishek Kadian*, Oleksandr Maksymets*, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, D. Parikh, and D. Batra, “Habitat: A Platform for Embodied AI Research,” in *IEEE ICCV*, 2019.
- [19] E. Kolve *et al.*, “Ai2-thor: An interactive 3d environment for visual ai,” *arxiv:1712.05474*, 2017.
- [20] X. Puig *et al.*, “Virtualhome: Simulating household activities via programs,” in *IEEE CVPR*, 2018.
- [21] A. X. Chang *et al.*, “Shapenet: An information-rich 3d model repository,” *arxiv:1512.03012*, 2015.
- [22] X. Wang, B. Zhou, Y. Shi, X. Chen, Q. Zhao, and K. Xu, “Shape2motion: Joint analysis of motion parts and attributes from 3d shapes,” in *IEEE CVPR*, 2019.
- [23] A. Kalervo, J. Ylioinas, M. Häikiö, A. Karhu, and J. Kannala, “Cubicasa5k: A dataset and an improved multi-task model for floorplan image analysis,” in *Scandinavian Conference on Image Analysis*, 2019.
- [24] H. Fu *et al.*, “3d-front: 3d furnished rooms with layouts and semantics,” *arxiv:2011.09127*, 2020.
- [25] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese, “Gibson env: Real-world perception for embodied agents,” in *CVPR*, 2018.
- [26] C. Gan *et al.*, “Threedworld: A platform for interactive multi-modal physical simulation,” *arxiv:2007.04954*, 2020.
- [27] E. Coumans *et al.*, “Bullet physics library,” *Open source: bulletphysics.org*, 2013.
- [28] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *IROS*, 2012.
- [29] J. Lee, M. X. Grey, S. Ha, T. Kunz, S. Jain, Y. Ye, S. S. Srinivasa, M. Stilman, and C. K. Liu, “Dart: Dynamic animation and robotics toolkit,” *Journal of Open Source Software*, 2018.
- [30] “Unity,” <https://unity.com/>, accessed: 2020-10-30.
- [31] J. Liang, V. Makovychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, “Gpu-accelerated robotic simulation for distributed reinforcement learning,” *CoRL*, 2018.
- [32] “Ode,” <https://ode.org/>, accessed: 2020-10-30.
- [33] Z. Erickson *et al.*, “Assistive gym: A physics simulation framework for assistive robotics,” in *IEEE ICRA*, 2020.
- [34] A. Chang *et al.*, “Matterport3d: Learning from rgb-d data in indoor environments,” in *3DV*, 2017.
- [35] F. Xia *et al.*, “Interactive gibson benchmark: A benchmark for interactive navigation in cluttered environments,” *IEEE RA-L*, 2020.
- [36] G. Brockman *et al.*, “Openai gym,” *arxiv:1606.01540*, 2016.
- [37] D. Merkel, “Docker: lightweight linux containers for consistent development and deployment,” *Linux journal*, 2014.
- [38] X. Meng, N. Ratliff, Y. Xiang, and D. Fox, “Scaling local control to large-scale topological navigation,” in *IEEE ICRA*, 2020.
- [39] —, “Neural autonomous navigation with riemannian motion policy,” in *IEEE ICRA*, 2019.
- [40] K. Kang *et al.*, “Generalization through simulation: Integrating simulated and real data into deep reinforcement learning for vision-based autonomous flight,” *IEEE ICRA*, 2019.
- [41] N. Hirose, F. Xia, R. Martín-Martín, A. Sadeghian, and S. Savarese, “Deep visual mpc-policy learning for navigation,” *IEEE RA-L*, 2019.
- [42] “iGibson Sim2real Challenge at CVPR2020.” [Online]. Available: <http://svl.stanford.edu/igibson/challenge2020.html>
- [43] “iGibson Challenge at CVPR2021.” [Online]. Available: <http://svl.stanford.edu/igibson/challenge.html>
- [44] C. Pérez-D’Arpino, C. Liu, P. Goebel, R. Martín-Martín, and S. Savarese, “Robot navigation in constrained pedestrian environments using reinforcement learning,” *IEEE ICRA*, 2021.
- [45] K. Mo, S. Zhu, A. X. Chang, L. Yi, S. Tripathi, L. J. Guibas, and H. Su, “Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding,” in *IEEE CVPR*, 2019.
- [46] C. Schlick, “An inexpensive brdf model for physically-based rendering,” in *Computer graphics forum*, 1994.
- [47] “Nadrin/pbr,” <https://github.com/Nadrin/PBR>, accessed: 2020-10-30.
- [48] L. Porzi, S. R. Bulo, A. Penate-Sanchez, E. Ricci, and F. Moreno-Noguer, “Learning depth-aware deep representations for robotic perception,” *IEEE RA-L*, 2016.
- [49] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova, “Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos,” in *AAAI*, 2019.
- [50] C. Xie, Y. Xiang, A. Mousavian, and D. Fox, “The best of both modes: Separately leveraging rgb and depth for unseen object instance segmentation,” in *CoRL*, 2020.
- [51] I. Nematollahi, O. Mees, L. Hermann, and W. Burgard, “Hindsight for foresight: Unsupervised structured dynamics models from physical interaction,” *IEEE IROS*, 2020.
- [52] C. Choi and H. I. Christensen, “Rgb-d object pose estimation in unstructured environments,” *Robotics and Autonomous Systems*, 2016.
- [53] C. Wang *et al.*, “6-pack: Category-level 6d pose tracker with anchor-based keypoints,” in *IEEE ICRA*, 2020.
- [54] Y. Hu, J. Hugonot, P. Fua, and M. Salzmann, “Segmentation-driven 6d object pose estimation,” in *IEEE CVPR*, 2019.
- [55] M. Yan, Q. Sun, I. Frosio, S. Tyree, and J. Kautz, “How to close sim-real gap? transfer with segmentation!” *arxiv:2005.07695*, 2020.
- [56] F. Sadeghi and S. Levine, “Cad2rl: Real single-image flight without a single real image,” *RSS*, 2016.
- [57] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, “Domain randomization for transferring deep neural networks from simulation to the real world,” in *IROS*, 2017.
- [58] S. James *et al.*, “Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks,” in *IEEE CVPR*, 2019.
- [59] M. Andrychowicz *et al.*, “Learning dexterous in-hand manipulation,” *IJRR*, 2020.
- [60] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [61] S. M. LaValle, “Rapidly-exploring random trees: A new tool for path planning,” Department of Computer Science; Iowa State University, Tech. Rep., 1998.
- [62] A. H. Qureshi and Y. Ayaz, “Intelligent bidirectional rapidly-exploring random trees for optimal motion planning in complex cluttered environments,” *Robotics and Autonomous Systems*, 2015.
- [63] R. Bohlin and L. E. Kavraki, “Path planning using lazy prm,” in *IEEE ICRA*, 2000.
- [64] Caelan Reed Garrett, “PyBullet Planning.” <https://pypi.org/project/pybullet-planning/>, 2018.
- [65] K. Hauser and V. Ng-Thow-Hing, “Fast smoothing of manipulator trajectories using optimal bounded-acceleration shortcuts,” in *IEEE ICRA*, 2010.
- [66] P. Anderson *et al.*, “On evaluation of embodied navigation agents,” *arxiv:1807.06757*, 2018.
- [67] A. Mandlekar *et al.*, “Roboturk: A crowdsourcing platform for robotic skill learning through imitation,” in *CoRL*, 2018.
- [68] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on MICCAI*, 2015.
- [69] H. Van Hasselt, A. Guez, and D. Silver, “Deep reinforcement learning with double q-learning,” in *AAAI*, 2016.