

ROS-X-Habitat: Bridging the ROS Ecosystem with Embodied AI

Guanxiong Chen, Haoyu Yang, Ian M. Mitchell
Department of Computer Science
The University of British Columbia
Vancouver, BC
Email: {gxchen, yhymanon, mitchell}@cs.ubc.ca

Abstract—We introduce ROS-X-Habitat, a software interface that bridges the AI Habitat platform for embodied learning-based agents with other robotics resources via ROS. This interface not only offers standardized communication protocols between embodied agents and simulators, but also enables physically and photorealistic simulation that benefits the training and/or testing of vision-based embodied agents. With this interface, roboticists can evaluate their own Habitat RL agents in another ROS-based simulator or use Habitat Sim v2 as the test bed for their own robotic algorithms. Through in silico experiments, we demonstrate that ROS-X-Habitat has minimal impact on the navigation performance and simulation speed of a Habitat RGBD agent; that a standard set of ROS mapping, planning and navigation tools can run in Habitat Sim v2; and that a Habitat agent can run in the standard ROS simulator Gazebo.

Keywords—Embodied Visual Agents; Robotic Middleware; Simulation

I. INTRODUCTION

Since the earliest days of robotics, researchers have sought to build embodied agents to perform a variety of jobs, such as assistive tasks in factories [1] or wildfire surveillance [2]. Following tremendous advancements in deep learning and convolutional neural networks over the past decade, researchers have been able to develop learning-based embodied agents that interact with the real world on the basis of visual observations [3]. Software platforms such as OpenAI Gym [4], Unity ML-Agents Toolkit [5], Nvidia Isaac Sim [6] and AI Habitat [7] have emerged to address the community's need for training and evaluating learning-based embodied visual agents end-to-end. Our research group was particularly intrigued by the AI Habitat platform, because this open source platform offers developers not only flexibility to tailor the software to their own needs, but also a high-performance photorealistic simulator recently enhanced with the Bullet physics engine [8], as well as access to a sizeable library of visually-rich scanned 3D environments with which to bridge the *sim2real* performance gap for visual agents [9].

Even though these platforms allow roboticists to reuse existing learning algorithms and train visual agents in simulators with ease, there is a critical step to using them for embodied agents which is only partially addressed: Connecting the trained agent with a robot. After training a learning-based agent in simulation one would like to take advantage of the extensive set of tools and knowledge from the robotics

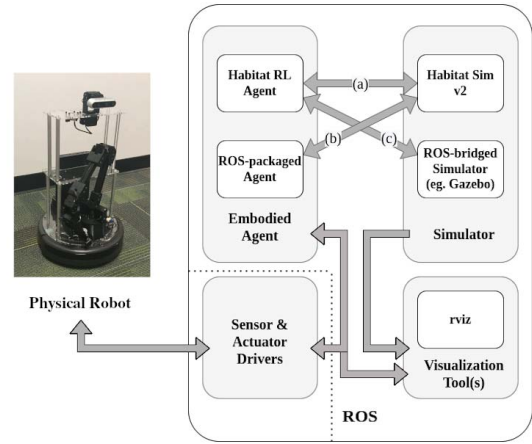


Fig. 1: High-level overview of `ros_x_habitat`'s architecture. In this paper we demonstrate the system's operation under modes (a), (b) and (c) to interface ROS or Habitat agents with a ROS-based simulator or Habitat Sim v2. Although not demonstrated in the paper, the modular design of `ros_x_habitat` allows easy embodiment of a Habitat visual navigation agent on a physical robot (such as the LoCoBot [12]).

community to make it easy to embody that agent. One particularly popular tool from the robotics community is ROS, a robotics-focused middleware platform with extensive support for classical robotic mapping, planning and control algorithms, as well as drivers for a wide variety of computing, sensing and actuation hardware. But ROS' support for direct interfacing with RL agents is limited, and Gazebo—the standard simulation environment used for ROS systems—suffers from two significant shortcomings when applied to this task: (i) It lacks the level of photorealism required to allow vision-based agents trained with synthetic inputs to perform well in real world scenarios [9], [10]; (ii) It cannot match the simulation speed of tools specifically designed to train large-scale reinforcement learning (RL) agents [11].

In order to take advantage of the strengths and overcome the weaknesses of these two independent sets of tools, in this paper we present ROS-X-Habitat (“ROS-Cross-Habitat”), an interface that bridges the AI Habitat training platform with the ROS ecosystem. Figure 1 shows a simplified view of

the interface’s architecture. The interface makes the following contributions to the robotics community: (i) It allows embodied agents to leverage Habitat Sim v2’s physics-based simulation capability through the interface. We demonstrate that the performance of a Habitat agent trained in a simulation environment without physics degrades only slightly when physics is turned on via the interface. This result suggests that high throughput training *without* the overhead of ROS and physics can produce agents that are still effective once these additional layers are added. (ii) It allows AI Habitat’s vision-based RL agents to be evaluated in ROS, so RL agent developers can take advantage of ROS’s rich set of tools and community support, as well as Gazebo’s ability to let users author 3D assets and customary testing scenes. Although not demonstrated here, this bridge to ROS dramatically shortens the path to embodying the Habitat agents in a physical robot. (iii) It allows other embodied agents implemented in ROS packages—for example, classical navigation algorithms such as [13], [14]—access to the rich 3D environments available in Habitat Sim v2. Consequently, algorithms from other sources can be evaluated in these photorealistic environments. Additional details and features of the project can be found in [15], [16]. The source code of the interface is available at https://github.com/ericchen321/ros_x_habitat.

II. RELATED WORK

Based on the objectives of our work, we review previous work in three relevant areas of research: (i) embodied agents, (ii) robotic middleware, (iii) simulators and datasets for training and evaluating learning-based embodied agents or for evaluating classical embodied agents in general.

A. Embodied Agents

We consider two categories of embodied agents commonly used to complete navigation tasks: Classical robotics approaches and learning-based.

Most commonly deployed classical embodied agents do navigation in two phases: Construct a map of the environment using, for example, a SLAM algorithm (such as [17], [18], [19]), and then use the map to plan out a path to the goal position (such as [13], [14]). While many packages that follow this approach are available in ROS, most recent photorealistic, physics-capable simulators (see Section II-B) do not interface to ROS and so evaluating classical agents in those simulators is difficult. So-called “end-to-end” learning-based agents use a neural network to produce a sequence of actions directly from visual observations and localization data without relying on prior maps [7], [20], [21]. Another popular approach is to combine learning-based agents with classical mapping-then-planning [22], [23], [24]. But none of these frameworks connect directly to ROS.

Although not fundamental to their designs, a common distinction between classical and visual learning-based agents is that the former operate in a continuous action space while the latter (especially RL agents) are often trained to produce

discrete actions. Continuous action spaces are more representative of how physical robots actuate [25], but training for these spaces has higher computational cost [26]. As an example, Habitat’s default PointGoal navigation agents have an action space consisting of four actions [7]: `move_forward`, `turn_left`, `turn_right`, and `stop`. To simulate these actions, Habitat Sim teleports the robot from one state to another without taking account of interactions between the agent and other objects at intermediate states. It is certainly possible to map from discrete to continuous actions, but it is not clear a priori that embodying a discrete agent into a continuous action space in this manner will produce good results.

B. Simulators and Datasets

The exploding interest in RL agents, and particularly vision-based RL agents, has put a premium on simulation speed and photorealism. Gazebo [27] has been the canonical choice for simulation in the ROS community; it benefits from extensive community support and a huge set of community-shared prebuilt assets, but lacks photorealism and high simulation speed. Unity is a powerful game engine, and the Unity ML-Agents Toolkit available to researchers provides simulation environments suitable for embodied agents [5], but the toolkit does not provide photorealistic simulation spaces by default and lacks compatibility with off-the-shelf 3D datasets such as Replica [28] or Matterport3D [29]. Isaac Sim [6] shows great promise in terms of configurability and photorealism, but is not currently open-source. The recently released Sapien [30] platform offers a photorealistic and physics-rich simulated environment, but currently provides limited support for tasks other than motion planning.

In this paper we explore the use of Habitat Sim v2 from the AI Habitat platform [31] for several reasons: (i) Extremely high speed simulation. This feature is particularly useful for RL agents, since agent performance may continue to improve even after many millions of training steps. (ii) Photorealistic rendering of scanned spaces. Habitat Sim v2 can render photorealistic scenes (including depth maps) from Habitat’s native datasets such as Replica [28] and ReplicaCAD [31]. The Habitat framework’s modular design [7], [31] offers a seamless interface to many prior photorealistic 3D scene datasets, including Matterport3D [29] and Gibson [32]. For this paper we use the Matterport3D test set, since it is the only publicly available dataset for which Habitat navigation results have been reported against which we can compare (albeit for v1 agents) [7]. (iii) Simulation of many different tasks. We focus here on PointGoal navigation [7] from Habitat v1, but object picking has also been demonstrated [31].

C. Robotics Middleware

Robotics middleware [33] is an abstraction layer that resides between the robotics software and the operating system (which is itself abstracting the underlying hardware). Middleware provides standardized APIs to sensors, actuators, and communication; design modularity; and portability.

A variety of robotics middleware systems have been developed (for example, see [34]), each with its own strengths and limitations. For example, while Microsoft RDS supports multiple programming languages, it runs only on the Windows OS[35]. OROCOS offers its own optimized runtime environment for real-time applications, but it does not have a graphical environment for drag-and-drop development nor a simulation environment [35]. ROS (Robot Operating System) is a free, open-source and popular robotics middleware introduced in 2007 [36]. Among its features: (i) It promotes modular, robust and potentially distributed designs; (ii) The huge user community has generated thousands of ready-to-use packages, including drivers for all common robotics hardware [36], [37], [38].

Given the popularity of ROS, it should come as no surprise that others have sought to build an interface between an embodied AI platform and the ROS ecosystem. PathBench is a unified interface that allows a multitude of path-planning algorithms to be evaluated under ROS and Gazebo [39]. Zamora et al. [40] used ROS as a bridge between OpenAI Gym and Gazebo, but did not consider other simulation environments which could provide photorealism. The Habitat-PyRobot Bridge (HaPy) [41] is particularly related to our work: it provides integration between Habitat agents and PyRobot, which is itself an interface on top of ROS to control different robots [42]. A key distinction between our work and HaPy-PyRobot is a design decision about how much the user is exposed to ROS. PyRobot abstracts ROS components away from top-level APIs [42], thus requiring less familiarity with ROS from its users. ROS-X-Habitat exposes ROS components to a greater extent, thus requiring more familiarity with ROS but allowing users greater flexibility to take advantage of the full ROS ecosystem. Also, HaPy does not currently support physics-based simulation from Habitat Sim.

III. SYSTEM DESIGN

A. Requirements

The interface should allow users to:

- 1) **Leverage Habitat Sim v2's full physics-based simulation capability to provide a more realistic environment for embodied agents.** This step is non-trivial because (at the current time) Habitat Sim v2 cannot simulate agents which produce discrete actions with full dynamic physics. In order to do so we have implemented API calls in the interface to enable physics-based simulation. Section IV demonstrates success on this requirement for a learned Habitat RGBD navigation agent, and Section V demonstrates success for a classical ROS planner.
- 2) **Evaluate an agent inside a simulator through the interface with minimal performance impact.** Section IV reports a number of experiments designed to gauge the effect of the ROS middleware on the performance of the Habitat platform components.
- 3) **Deploy ROS-based planners to navigate within Habitat Sim v2, thus allowing researchers to test planners on a huge variety of photorealistic 3D scenes.** Section V

demonstrates that this combination is feasible using the Matterport 3D dataset.

- 4) **Evaluate Habitat agents in other simulation environments with ROS bridges or on physical hardware.** Section VI demonstrates this capability in Gazebo, and once we know that this combination works the jump to other simulators or a real robot is much less daunting.
- 5) **Extend or customize its functionality for bridging with other tools from both the ROS and Habitat frameworks.** In Section III-B we show the modular design of ROS-X-Habitat. The exposure of ROS nodes, topics and services allows users to interface to other ROS tools (such as RViz) or Habitat's visualization modules.

B. Architecture

We designed ROS-X-Habitat as a ROS Noetic package (`ros_x_habitat`), written in Python and consisting of a collection of ROS nodes, topics and services. Following the five requirements we established, we defined three operating modes for the interface, as illustrated in Figure 2.

- (a) **A Habitat agent navigates in a Habitat Sim-rendered scene (with or without physics) through ROS.** Note that this is not an expected use case for `ros_x_habitat`; after all, running a Habitat agent in Habitat Sim is precisely what AI Habitat is already designed to do. We include this operating mode to demonstrate that (i) a learning-based embodied agent trained in an environment without physics can be simulated with physics, and (ii) the `ros_x_habitat` interface does not introduce any performance penalty and the runtime overhead is not excessive.
- (b) **A ROS-based planner (`move_base` [14] in our case) navigates in a Habitat Sim-rendered scene (with physics).** Here we replace `/agent_node` with a ROS-based planner node in the control loop. During navigation, we feed visual observations and a prebuilt 2D occupancy map to the ROS planner for localization and trajectory planning. Note that (i) depth observations are emulated as laser readings with `/depthimage_to_laserscan` [43] before being injected into the planner.
- (c) **A Habitat agent navigates in a Gazebo-rendered scene.** Although not demonstrated here, the nodes we designed for this mode also make it easy to swap in other ROS-bridged simulators or physical robots because those entities typically expose a ROS interface similar to Gazebo's.

IV. ANALYSIS OF IMPACT OF INTERFACE AND PHYSICS

A ROS bridge for the Habitat platform would be ineffective if it cannot leverage Habitat Sim v2's physics-based simulation. While Habitat Sim v2 offers physics-based simulation, the vision-based navigation agents that come with the platform were trained without physics, and Habitat Sim v2 currently cannot simulate these agents accurately with physics because they produce discrete actions. Moreover, if upon introducing physics the bridge either significantly degrades the navigation performance of an embodied agent or leads to unacceptable

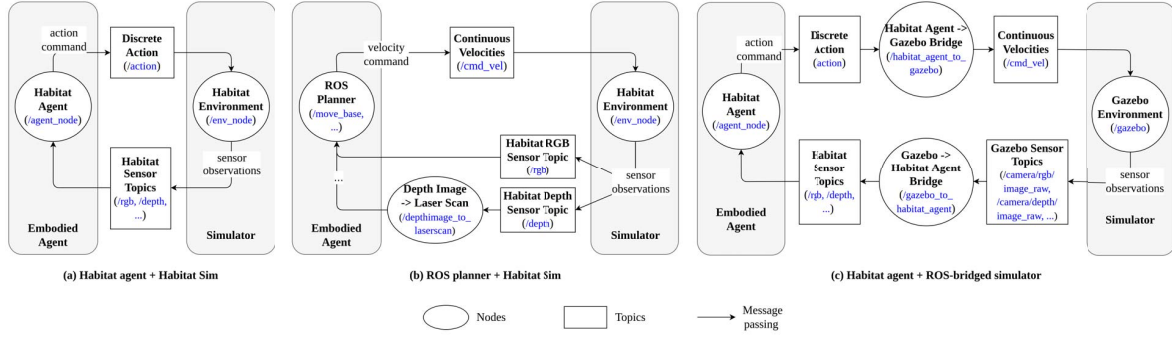


Fig. 2: The `ros_x_habitat` architecture under each operating mode. (a) Simulating a Habitat RL agent in Habitat Sim v2. If physics is turned on, each discrete action would be converted to a sequence of velocity commands in the node encapsulating the Habitat environment (`env_node`) following Algorithm 1. (b) Simulating a ROS-based planner in a physics-enabled and photorealistic Habitat Sim v2 environment. (c) Simulating a Habitat RL agent in a Gazebo environment. Note that for simplicity, we omitted nodes, topics and services not being the gist of the interface’s operation—for example, map topics and nodes under operating mode (b).

TABLE I: Simulation in Habitat Sim v2 with versus without physics. Habitat Sim v2 does not currently support simulation with physics for its visual navigation agents, but `ros_x_habitat` allows such cross-simulation.

	Agent’s Geometry	Agent’s Physical Properties	Simulation of Motion
Simulation without physics	A cylinder 0.1m in radius, 1.5m in height.	Friction coefficient is undefined; mass defined but not used for simulation.	World state advances by one <i>action step</i> for each discrete action simulated; for each action the agent is translated or rotated instantaneously; no forces simulated.
Simulation with physics	Defined by the 3D asset attached to agent’s scene node.	Friction coefficient is defined; mass is defined and used to compute dynamics.	World state advances by one <i>continuous step</i> for each velocity command; forces fully simulated.

runtime overhead it would still be an impractical design. We dedicate this section to answer two research questions inspired by these issues.

(RQ1) Given a Habitat navigation agent that was trained in Habitat Sim using discrete time and without physics, how effective is its navigation in the same environments but with continuous time and physics enabled? Having a Habitat agent’s discrete actions converted to a sequence of velocity commands allows us to leverage Habitat Sim’s physics engine to simulate the actuation process in a more realistic fashion. We would like to measure how much of an impact physics has on navigation performance and execution speed.

(RQ2) Does our implementation using ROS middleware impair navigation performance or introduce unacceptable runtime overhead? First, we would like to verify that adding the ROS interface does not alter a Habitat agent’s navigation performance within Habitat Sim. Second, we would like to measure how much ROS overhead impacts execution speed. If the simulation throughput drops below real-time—one simulated second takes more than one second to simulate computationally—then we lose a significant benefit of simulation compared to testing in the real-world.

Think of Operating Mode (a) of `ros_x_habitat` as enabling a form of regression test: The goal is to demonstrate that physics and/or the ROS interface do not break the impressive

capabilities of pre-trained Habitat agents. The research questions specify our regression criteria, and the subsections below describe the experiments used to test them (Section IV-A) and the results of those experiments (Section IV-B).

A. Experiment Setup

Agent, task and dataset. We evaluate the Habitat v2 RGBD agent [31] on the PointGoal navigation task (in which the agent navigates from a pre-set initial position to a goal position [7]). The agent was trained in Habitat Sim v2, but without physics and in discrete time. It can output one of four discrete action commands at each timestep: `move_forward`, `turn_left`, `turn_right` or `stop` [7]. We used the 1,008 navigation episodes (each episode is an instance of the task) from the Matterport3D test set [29]. Because the original Habitat RGBD agent was evaluated on the Matterport3D dataset, using it leads to more meaningful benchmarking.

Evaluation metrics. We employ the Success Weighted by Path Length (`spl`) metric to evaluate the RGBD agent’s navigation performance in each episode [7], [31]. This metric lies in the range $[0, 1]$ and measures the length of the traversed path relative to the shortest path from the source to the destination; closer to 1 implies a shorter path and thus better performance. To capture our interface’s impact on execution speed we measured the total wall clock runtime

(*running_time*) of evaluating an agent’s navigation over all 1,008 episodes, which accounts for overheads introduced by ROS, such as initialization and inter-process communication.

Enabling physics and mapping from discrete to continuous action space. In Table I we summarize how Habitat Sim v2 operates with and without physics, where (i) The 3D asset we attached to our agent is the *LoCoBot* model provided in Habitat Sim’s code base [7]; (ii) An *action step* is defined as one update of the world’s state due to the agent completing an action in a discrete time simulation without physics; (iii) A *continuous step* is defined as the advancement of a world’s state over a predefined time interval ($\frac{1}{60}$ second in our experiments) in a continuous time simulation with physics enabled. As suggested in Table I, simulation without physics teleports the agent instantaneously to the destination position of the discrete action, whereas simulation with physics requires the agent to move through space at a specified velocity. To transfer the existing RGBD agent, or any agent that outputs discrete actions, to the continuous time simulation with physics, we use Algorithm 1 to map each discrete action to a sequence of velocity commands, where (i) *control_period* is a user-supplied parameter and defines the time in seconds it takes an agent to complete an action step. Set to 1 second in our experiments; (ii) *steps_per_sec* defines the number of continuous steps Habitat Sim advances for each second of simulated time. Set to 60 in our experiments; (iii) the angular velocities are measured in degrees / second; and (iv) the linear velocities are measured in meters / second.

Algorithm 1 Convert a discrete action from a Habitat agent to a sequence of velocities. The conversion allows the interface to invoke Habitat Sim v2’s built-in API (*step_physics()*) to simulate the action with physics.

Require: a discrete *action* as one of the following:
 move_forward, *turn_left*, or *turn_right*
 1: initialize *linear_velocity*, *angular_velocity* as zero vectors
 2: *num_steps* = *control_period* · *steps_per_sec*
 3: **if** *action* == *move_forward* **then**
 4: *linear_velocity* = [0.25/*control_period*, 0, 0]
 5: **else if** *action* == *turn_left* **then**
 6: *angular_velocity* = [0, 0, 10.0/*control_period*]
 7: **else if** *action* == *turn_right* **then**
 8: *angular_velocity* = [0, 0, -10.0/*control_period*]
 9: **end if**
 10: **for** *count_steps* = 1, 2, ..., *num_steps* **do**
 11: *env.step_physics*(*linear_velocity*, *angular_velocity*)
 12: **end for**

Experiment Configurations. We conducted our experiments on four configurations in order to independently observe the impact of introducing physics-based simulation and adding ROS. (i) *-Physics & -ROS*. We have the Habitat RGBD agent actuate in its default, discrete action space without using the ROS middleware, and we run Habitat Sim without




Fig. 3: Distribution of *spl* under the four experimental configurations. Failed episodes have *spl* at 0. Under Configurations (a) and (b), the mean *spl* produced by the v2 RGBD agent is 0.495 without physics, which is close to the mean *spl* (0.53) reported for the v1 RGBD agent in [7]. Under Configurations (c) and (d), the *spl* dropped slightly to 0.455 after physics was enabled.

physics turned on. This setting is the configuration in which the agent was trained, and serves as a baseline for our experiment. (ii) *+Physics & -ROS*. We enable physics-based simulation, but do not use ROS. (iii) *-Physics & +ROS*. The Habitat RGBD agent communicates with Habitat Sim v2 through the ROS interface, as shown in Figure 2(a). The agent still navigates using its discrete action space, and the simulation is run without physics. (iv) *+Physics & +ROS*. The combination of the previous two settings.

Evaluation platform. We ran all experiments on a desktop with an i7-10700K CPU, 64 GB of RAM, and an NVIDIA RTX 3070 GPU under Ubuntu 20.04, ROS Noetic and Habitat v0.2.0.

B. Results and Analysis

In this section we present the results from our experiments and seek to answer the two research questions posed earlier.

(RQ1) Given a Habitat navigation agent that was trained in Habitat Sim using discrete time and without physics, how effective is its navigation in the same environments but with continuous time and physics enabled? Comparing subplots (a) and (c) in Figure 3 we can see the effect that enabling physics has on the distribution of *spl* over the 1,008 episodes. We see that the number of failed episodes (those with *spl* = 0.0) increased (by 129) after we mapped actions to velocities and introduced physics-based simulation, but the average *spl* only dropped slightly (from 0.495 to 0.455). We visually examined the 129 episodes in which the agent navigated successfully in the discrete action space but failed in the continuous action space. For the majority of them (80/129) the failure was due to the robot being “stuck”, i.e. the agent is not able to move forward but remains within

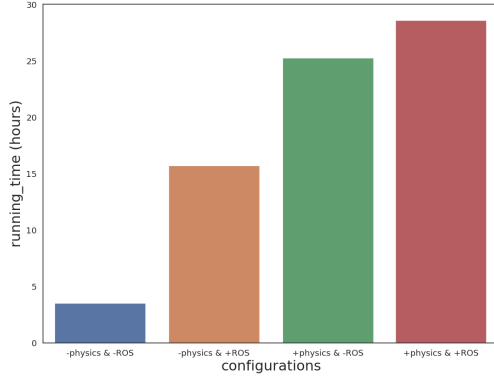


Fig. 4: Total runtime under the four experimental configurations for all 1,008 episodes.

the scene. In some episodes the agent was stuck because the agent’s movement is constrained to two dimensions by our implementation of Algorithm 1: It was not able to climb stairs when `move_forward` is converted to a sequence of planar velocities. This problem does not manifest in the discrete time mode without physics: it appears that the agent is automatically teleported to the correct height for whatever planar position it should occupy.

Figure 4 shows our timing result. Adding just physics (without ROS) slows Habitat by a factor of roughly eight, as extra time was required for simulating a discrete action as a sequence of continuous actions (Algorithm 1), and for simulator reset (to load and delete physics-based object assets and reconfigure the simulator).

We conclude that using physics will significantly reduce the throughput of the simulation engine, but the modest performance degradation confirms that training RL agents in the simplistic but high-throughput discrete action configuration without physics can produce reasonable results even for agents intended to be used in more realistic environments [26], and the full physics need not be turned on until validation or even final testing.

(RQ2) Does our implementation using ROS middleware impair navigation performance or introduce unacceptable runtime overhead? In terms of navigation performance, once we removed non-determinism from the system we were able to show that the introduction of ROS did not affect navigation performance regardless of whether physics was enabled: Each episode had the exact same number of steps with and without ROS, and the SPL discrepancies were at the level of floating point round-off. Comparing the distributions in the left and right columns of figure 3 shows the latter effect qualitatively.

Turning to execution time: Figure 4 shows that adding ROS increases total runtime by roughly a factor of five (without physics) and less than 20% (with physics). The increases are due to the overheads from inter-process communication between ROS nodes, as well as intra-process communication




Fig. 5: ROS-based planner `move_base` navigating in Matterport3D scene `2t7WUuJeko7` simulated by Habitat Sim v2. (a) The agent’s final position overlayed on top of the map build by `rtab_map_ros`. (b) Top-down map of the scene from Habitat Sim v2. The blue curve indicates the agent’s trajectory. Blue square: starting position of the agent. Yellow dot: final position. Red dot: goal position.




Fig. 6: ROS-based planner `move_base` navigating in Matterport3D scene `RPmz2sHmrrY` simulated by Habitat Sim. (a) The agent’s final position and laser-scanned map. (b) Top-down map from Habitat Sim.

between threads running various service handlers and subscriber callbacks within each node. Although the increase in running time is significant, it is acceptable: with overhead from both ROS and physics taken into account, simulating over 274,000 action steps across the full set of episodes in less than 30 hours implies an average computation time of less than 0.4 seconds for each action step lasting one second of simulated time; consequently, the throughput is more than twice real-time.

V. NAVIGATION OF ROS-BASED PLANNERS IN HABITAT SIM

For robotics researchers interested in designing and testing navigation algorithms—be they classical or RL—we believe that this configuration of `ros_x_habitat` will be most useful, as Habitat Sim provides a novel, high throughput, photorealistic simulation environment in which to test such algorithms. In this section we do not seek to show the merits of a particular mapping, planning or navigation system, but simply that some standard packages from ROS can be easily and successfully connected to Habitat Sim in the configuration shown in Figure 2(b).




Fig. 7: Maps showing the Habitat v2 RGBD agent navigating in three episodes in the House scene. The blue spheres indicate the starting positions of each episode; the red spheres indicate the goal positions of each episode. The green curves represent paths traversed by the agent.

First, we mapped two scenes from the Matterport3D dataset with `rtab_map_ros` [17]. Second, we attached the `LoCoBot` [7] asset to the agent in Habitat Sim (the same asset used in the previous section’s experiments). Finally, we manually set a goal position for the `move_base` [14] planner. Figures 5 and 6 show the generated map (left) and ground-truth map (right) overlaid with the final agent position and sensor readings (left) and path (right) for the two scenes. The planner failed to reach the goal in Figure 5 but succeeded in Figure 6.

We observed during the simulations that the agent often had a hard time localizing itself, especially in cluttered regions. We expect that with some tuning of the mapping, planning and navigation algorithms’ parameters we could achieve better performance.

VI. NAVIGATION OF HABITAT AGENTS IN GAZEBO

For RL researchers interested in testing their agents on real robots, this configuration demonstrates the benefits of `ros_x_habitat`. Researchers familiar with ROS know that despite the significant `sim2real` gap in the Gazebo simulator, it can be an effective first target during design and testing because it will expose whole classes of common design bugs, such as incorrectly typed or connected data flows, coordinate transform errors, and gross timing issues.

Although we do not demonstrate a Habitat agent connected to a physical robot in this paper, by demonstrating a Habitat agent connected to Gazebo we show that `ros_x_habitat` can allow connection of Habitat agents to other simulation environments that have ROS bridges. Note that we do not expect the Habitat agent to perform particularly well in this Gazebo environment since it was trained in the much richer visual and geometric environments of Matterport3D in Habitat Sim.

Injecting the Matterport3D dataset into Gazebo is not practical, so we used the House scene from the `turtlebot3_gazebo` package [44] to build three navigation episodes of increasing path length as shown in Figure 7. Then we instantiate a Habitat v2 RGBD agent inside each episode, and let it navigate until it reaches the goal or has generated 500 actions. Since we only intend to show that this configuration works, we do not report quantitative metrics such

as `spl` for these episodes. We see that the agent succeeded in the two shorter of the three episodes, although in the middle episode the agent took a path much longer than would be optimal.

VII. CONCLUSION AND FUTURE WORK

We presented an extensible interface that bridges the AI Habitat platform with ROS middleware. Through this interface, researchers can develop and test their embodied agents in both established simulators such as Gazebo or the state-of-the-art Habitat Sim v2. Moreover, our interface implements the necessary transformations to allow agents trained in the original, high-throughput discrete time Habitat Sim environment without physics to navigate in the more realistic continuous time environment with physics available in Habitat Sim v2. Using the SPL metric, we show that our interface introduces negligible impacts on the navigation performance of embodied agents in a discrete action space, and we show that if the agents are actuated in a continuous action space with physics-based simulation enabled, the run time overhead is still acceptable. Although we only include two simulator options in our current work, our modular design allows easy integration of additional robotic embodiments, such as other simulators or physical robots.

The interface is still in its infancy. Some planned work for the future of `ros_x_habitat` include: (i) Allowing the simulation of embodied agents trained for tasks other than navigation, such as the object picking task introduced in Habitat v2 [31]. (ii) Allowing classical ROS planners to be evaluated on Habitat’s navigation metrics, such as `spl` when simulated under Habitat Sim v2. (iii) Deploying an embodied Habitat agent on physical robot(s).

ACKNOWLEDGMENT

The authors were partially supported by National Science and Engineering Research Council of Canada (NSERC) Discovery Grant #RGPIN-2017-04543. The authors would also like to thank Bruce Cui from the Department of Mechanical Engineering at UBC for his initial work on `ros_x_habitat`, and the AI Habitat team from Facebook AI Research, including Prof. Dhruv Batra, Alex Clegg, Prof. Manolis Savva, and Erik Wijmans for their incredible support throughout our development of the interface.

REFERENCES

- [1] H. Oliff, Y. Liu, M. Kumar, M. Williams, and M. Ryan, “Reinforcement learning for facilitating human-robot-interaction in manufacturing,” *Journal of Manufacturing Systems*, vol. 56, pp. 326–340, 2020.
- [2] K. D. Julian and M. J. Kochenderfer, “Distributed wildfire surveillance with autonomous aircraft using deep reinforcement learning,” *Journal of Guidance, Control, and Dynamics*, vol. 42, no. 8, pp. 1768–1778, 2019.
- [3] J. Duan, S. Yu, H. L. Tan, H. Zhu, and C. Tan, “A survey of embodied ai: From simulators to research tasks,” *IEEE Transactions on Emerging Topics in Computational Intelligence*, pp. 1–15, 2022.
- [4] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.

- [5] A. Juliani, V. Berges, E. Vckay, Y. Gao, H. Henry, M. Mattar, and D. Lange, "Unity: A general platform for intelligent agents," *CoRR*, vol. abs/1809.02627, 2018. [Online]. Available: <http://arxiv.org/abs/1809.02627>
- [6] F. Monteiro, A. L. Vieira e Silva, J. M. Teixeira, and V. Teichrieb, "Simulating real robots in virtual environments using nvidia's isaac sdk," 10 2019, pp. 47–48.
- [7] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied ai research," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9339–9347.
- [8] "Bullet physics sdk," <https://github.com/bulletphysics/bullet3>, accessed: 2021-09-03.
- [9] A. Garcia-Garcia, P. Martinez-Gonzalez, S. Oprea, J. A. Castro-Vargas, S. Orts-Escolano, J. Garcia-Rodriguez, and A. Jover-Alvarez, "The robotrix: An extremely photorealistic and very-large-scale indoor dataset of sequences with robot trajectories and interactions," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 6790–6797.
- [10] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig, "Virtual worlds as proxy for multi-object tracking analysis," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4340–4349.
- [11] Z. Liang, Z. Cai, M. Li, and W. Yang, "Parallel gym gazebo: a scalable parallel robot deep reinforcement learning platform," in *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*, 2019, pp. 206–213.
- [12] "Locobot an open source low cost robot," <http://www.locobot.org/>, accessed: 2022-02-24.
- [13] "dwa_local_planner - ros wiki." [Online]. Available: http://wiki.ros.org/dwa_local_planner
- [14] "move_base - ros wiki." [Online]. Available: http://wiki.ros.org/move_base
- [15] G. Chen, H. Yang, and I. M. Mitchell, "Ros-x-habitat: Bridging the ros ecosystem with embodied ai," *arXiv preprint arXiv:2109.07703*, 2021.
- [16] H. Yang, "Ros-x-habitat: Bridging the ros ecosystem with embodied ai," 2022. [Online]. Available: <http://hdl.handle.net/2429/81124>
- [17] M. Labbé and F. Michaud, "Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation," *Journal of Field Robotics*, vol. 36, no. 2, pp. 416–446, 2019. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21831>
- [18] S. Kohlbrecher, J. Meyer, T. Graber, K. Petersen, U. Klingauf, and O. von Stryk, "Hector open source modules for autonomous mapping and navigation with rescue robots," in *RoboCup 2013: Robot World Cup XVII*, S. Behnke, M. Veloso, A. Visser, and R. Xiong, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 624–631.
- [19] "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [20] L. Tai, G. Paolo, and M. Liu, "Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 31–36.
- [21] S. Datta, O. Maksymets, J. Hoffman, S. Lee, D. Batra, and D. Parikh, "Integrating egocentric localization for more realistic point-goal navigation agents," *arXiv preprint arXiv:2009.03231*, 2020.
- [22] G. Brunner, O. Richter, Y. Wang, and R. Wattenhofer, "Teaching a machine to read maps with deep reinforcement learning," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018. [Online]. Available: <https://ojs.aaai.org/index.php/AAAI/article/view/11645>
- [23] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov, "Learning to explore using active neural SLAM," *CoRR*, vol. abs/2004.05155, 2020. [Online]. Available: <https://arxiv.org/abs/2004.05155>
- [24] P. Nikdel, M. Chen, and R. Vaughan, "Recognizing and tracking high-level, human-meaningful navigation features of occupancy grid maps," in *2020 17th Conference on Computer and Robot Vision (CRV)*. IEEE, 2020, pp. 62–69.
- [25] W. Masson, P. Ranchod, and G. Konidaris, "Reinforcement learning with parameterized actions," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [26] E. Marchesini and A. Farinelli, "Discrete deep reinforcement learning for mapless navigation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 688–10 694.
- [27] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 3, 2004, pp. 2149–2154 vol.3.
- [28] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma *et al.*, "The replica dataset: A digital replica of indoor spaces," *arXiv preprint arXiv:1906.05797*, 2019.
- [29] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *arXiv preprint arXiv:1709.06158*, 2017.
- [30] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 097–11 107.
- [31] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. Chaplot, O. Maksymets *et al.*, "Habitat 2.0: Training home assistants to rearrange their habitat," *arXiv preprint arXiv:2106.14405*, 2021.
- [32] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese, "Gibson env: Real-world perception for embodied agents," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9068–9079.
- [33] A. Elkady and T. Sobh, "Robotics middleware: A comprehensive literature survey and attribute-based bibliography," *Journal of Robotics*, vol. 2012, 05 2012.
- [34] J. Kramer and M. Scheutz, "Development environments for autonomous mobile robots: A survey," *Auton. Robots*, vol. 22, pp. 101–132, 02 2007.
- [35] A. Hentout, A. Maoudj, and B. Bouzouia, "A survey of development frameworks for robotics," in *2016 8th International Conference on Modelling, Identification and Control (ICMIC)*, 2016, pp. 67–72.
- [36] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Ng, "Ros: an open-source robot operating system," vol. 3, 01 2009.
- [37] A. Hentout, A. Maoudj, and B. Bouzouia, "A survey of development frameworks for robotics," in *2016 8th International Conference on Modelling, Identification and Control (ICMIC)*, 2016, pp. 67–72.
- [38] G. Magyar, P. Sincak, and Z. Krizsán, "Comparison study of robotic middleware for robotic applications," *Advances in Intelligent Systems and Computing*, vol. 316, pp. 121–128, 01 2015.
- [39] A.-I. Toma, H.-Y. Hsueh, H. A. Jaafar, R. Murai, P. H. Kelly, and S. Saeedi, "Pathbench: A benchmarking platform for classical and learned path planning algorithms," in *2021 18th Conference on Robots and Vision (CRV)*. IEEE, 2021, pp. 79–86.
- [40] I. Zamora, N. G. Lopez, V. M. Vilches, and A. H. Cordero, "Extending the openai gym for robotics: a toolkit for reinforcement learning using ROS and gazebo," *CoRR*, vol. abs/1608.05742, 2016. [Online]. Available: <http://arxiv.org/abs/1608.05742>
- [41] A. Kadian, J. Truong, A. Gokaslan, A. Clegg, E. Wijmans, S. Lee, M. Savva, S. Chernova, and D. Batra, "Sim2real predictivity: Does evaluation in simulation predict real-world performance," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 08 2020.
- [42] A. Murali, T. Chen, K. V. Alwala, D. Gandhi, L. Pinto, S. Gupta, and A. Gupta, "Pyrobot: An open-source robotics framework for research and benchmarking," *CoRR*, vol. abs/1906.08236, 2019. [Online]. Available: <http://arxiv.org/abs/1906.08236>
- [43] "depthimage_to_laserscan - ros wiki." [Online]. Available: http://wiki.ros.org/depthimage_to_laserscan
- [44] "turtlebot3_gazebo - ros wiki." [Online]. Available: http://wiki.ros.org/turtlebot3_gazebo