

NeRF2Real: Sim2real Transfer of Vision-guided Bipedal Motion Skills using Neural Radiance Fields

Arunkumar Byravan¹, Jan Humplik¹, Leonard Hasenclever¹, Arthur Brussee¹, Francesco Nori, Tuomas Haarnoja, Ben Moran, Steven Bohez, Fereshteh Sadeghi, Bojan Vujatovic and Nicolas Heess
DeepMind, ¹Equal Contributions




Fig. 1: Zero-shot sim2real transfer results of vision-based bipedal locomotion policies trained using reinforcement learning in two separate simulations created using our NeRF2Real setup. **Left:** time lapse of a transfer result on a navigation task with a comparison of the robot’s head-mounted camera views vs NeRF renderings: ‘Real’: views from the real-robot, i.e. evaluation inputs to the policy; ‘NeRF’: train time NeRF rendered images. **Right:** time lapse of a result on a task where the robot has to push a ball towards the target region in front of the red cones. The policy was trained in simulation by overlaying simple rendering of an orange ball on top of the scene’s NeRF rendering.

Abstract—We present a system for applying sim2real approaches to “in the wild” scenes with realistic visuals, and to policies which rely on active perception using RGB cameras. Given a short video of a static scene collected using a generic phone, we learn the scene’s contact geometry and a function for novel view synthesis using a Neural Radiance Field (NeRF). We augment the NeRF rendering of the static scene by overlaying the rendering of other dynamic objects (e.g. the robot’s own body, a ball). A simulation is then created using the rendering engine in a physics simulator which computes contact dynamics from the static scene geometry (estimated from the NeRF volume density) and the dynamic objects’ geometry and physical properties (assumed known). We demonstrate that we can use this simulation to learn vision-based whole body navigation and ball pushing policies for a 20 degree-of-freedom humanoid robot with an actuated head-mounted RGB camera, and we successfully transfer these policies to a real robot.

I. INTRODUCTION

Thanks to progress in large-scale deep reinforcement learning and scalable simulation infrastructure, training control policies in simulation and transferring them to real robots (sim2real) has become a popular paradigm in robotics [1]–[4]. This approach avoids many of the issues such as state estimation, safety, and data efficiency which make it challenging to learn directly on hardware. However, creating accurate and realistic simulations is time consuming. Therefore, for sim2real to live up to its full potential, we must make it easier to recreate real scenes in simulation while accurately modelling how robots sense and interact with the world.

Reducing the gap between simulation and the real world often involves the collection of small amounts of data fol-

lowed by manual tuning, the use of established system identification tools, or more recently by learning neural network models of parts of the system, e.g. [1]. It is especially difficult to accurately model the geometry and visual appearance of unstructured scenes which affect how the robot makes contact with the world and how it senses its surroundings e.g. when using a RGB camera. The need for modeling RGB cameras can partially be alleviated by using depth sensors or LiDARs which are easier to simulate and thus have a smaller sim2real gap, but such a compromise can restrict the set of tasks a robot can learn. Existing approaches to photorealistic scene reconstruction and rendering, e.g. those used for the creation of the datasets in [5]–[8], work poorly in outdoor scenes and use specialized 3d scanning setups which are not widely available, hence limiting their applicability.

In this paper we begin to address some of these challenges, and describe a system for the semi-automated generation of simulation models for visually complex scenes with highly realistic rendering of RGB camera views and accurate geometry, primarily using videos from commodity mobile cameras. To this end, we take advantage of recent advances in neural scene representations using Neural Radiance Fields (NeRF) [9], [10]. NeRFs are a fast developing class of scene representations that allow synthesizing novel photorealistic views from a sparse set of input views. Unlike prior work, NeRFs can be learned directly from videos or photographs from commodity mobile devices and admit access not just to a rendering function but also to the underlying scene geometry. They can be trained within minutes [11], [12], work in

both indoor and outdoor settings and scale well even to large scenes such as city blocks [13]. Together with extensions to handle dynamic scenes [14], deformable objects [15], and scene decompositions with novel re-combinations of objects [16], NeRFs can enable a general system for recreating the visuals of real-world scenes in simulation.

Our primary contribution is an approach for combining NeRF scene representations, specifically the rendering and static geometry, learned from short (5-6) minute videos of a scene, with a physics simulation of dynamic objects such as a robot and a ball whose physical and visual properties are assumed known (see Fig. 2). We present a semi-automated pipeline for setting up these simulations and demonstrate that they have high enough fidelity to enable simulation-to-reality transfer of vision-guided control policies. Specifically, we use a physically accurate simulation of a 20 degree-of-freedom Robotis OP3 humanoid robot together with the NeRF and end-to-end deep reinforcement learning to train vision-based whole-body navigation and ball pushing policies and we show a strong alignment between the performance of these policies in simulation and when transferred zero-shot to a real robot (see Fig. 1 for a visualization of our results). Videos and scene reconstruction baseline comparisons can be seen at <https://sites.google.com/view/nerf2real>.

II. RELATED WORK

A. Neural radiance fields

Neural Radiance Fields (NeRF) [9] have recently become popular as an implicit scene representation capable of synthesizing novel photorealistic views. NeRF and its variants can represent accurate scene geometry [17], capture large scenes [13] and can be trained from photo collections without accurate localization or specialized hardware [18], [19]. Recent work has also shown training and rendering can be extremely fast and efficient [11], [12], or even be real-time on commodity handheld devices [20].

NeRF in robotics: NeRFs have been used in robotics for pose estimation [21], representation learning [22], grasping [23] and dynamics model learning [24]. There is also closely related work on obstacle avoidance within simulated NeRF environments leveraging a traditional state estimation and planning pipeline [25], but transfer to the real-world was not explored. Unlike this work we use reinforcement learning to train a policy that tightly integrates perception and control for a bipedal robot and demonstrate transfer to the real world on both visual navigation and object interaction tasks.

B. Visual navigation with (visually) realistic simulators

There is a long line of work on modeling real-world indoor scenes including datasets such as Matterport3D [5], Gibson [6], Replica [7] and Habitat-Matterport3D [26]. Unlike these datasets, which were predominantly created using purpose-built scanning setups, often with access to depth or LiDAR, we use a small amount of video data from off-the-shelf mobile cameras to train a NeRF to represent our scene.

Visual navigation in simulation: Several simulation suites have been proposed for embodied visual navigation

tasks combining 3D simulators with the different 3D scene datasets mentioned previously, such as Habitat [8], [27], iGibson [28] and AI2/ROBO-THOR [29], [30]. These simulators have been used for learning visual navigation policies [31]–[33], solve object-based navigation [34], [35], also incorporating language commands [36]. These approaches primarily consider dynamically simple platforms (wheeled robots) and operate purely in simulation.

Sim-to-real for visual navigation: Several works have demonstrated that policies trained in these photorealistic simulators can be transferred to real-world robots. The majority focuses on wheeled-base robots [4], [37], but some recent work has extended this to quadrupeds [38], [39]. These approaches use RGB-D sensors and/or LiDAR, assume access to localization, and in case of the quadrupeds, work on top of existing low-level controllers; all these assumptions help reduce the sim2real gap and thereby result in good transfer performance. In contrast, we successfully transfer whole-body vision-based control policies for a bipedal robot using only an RGB camera (which has a high sim2real gap) without access to localization or low-level controllers.

C. Sim2real in robotics

Sim2real transfer has made it possible to use reinforcement learning to solve several challenging real-world control problems. Careful system identification and techniques such as domain randomization [40], domain adaptation [41] and real-to-sim [42] have helped to reduce the discrepancies between simulation and reality for the system dynamics and sensor model, enabling successes on tasks such as Rubik’s cube solving with a dexterous hand [3], grasping [43], stacking [44], autonomous flight [45], quadruped [46]–[49] and biped locomotion [50]–[52]. We rely on many of the lessons learned in these works, and propose a system for high-fidelity replication of real-world scenes in simulation with which we demonstrate successful zero-shot transfer of complex vision-based policies on a 20 degree-of-freedom humanoid robot.

III. INTEGRATING NERF WITH A PHYSICS SIMULATOR

Fig. 2 presents an overview of our approach for recreating a static scene in simulation, and its extension to scenes with simple dynamic objects. Our approach consists of 6 steps: video recording, localization, NeRF training, post-processing to extract a rendering function and a collision mesh, and combining these with a physics simulator to create the simulation (see Fig 3). We describe each step below.

A. Capturing a video of the real world scene

We capture a short ~ 5 –6 minute video of the scene using an off-the-shelf mobile camera (Google Pixel 6’s rear camera in this work). A human operator walks around the scene and captures it from different viewpoints while ensuring that the camera moves slowly and evenly to reduce motion blur and minimize drastic viewpoint changes. For consistent lighting we set the white balance and brightness to a fixed (arbitrary) value. We found that high-resolution ($\geq 1080p$) videos led to better localization and improved NeRF results.

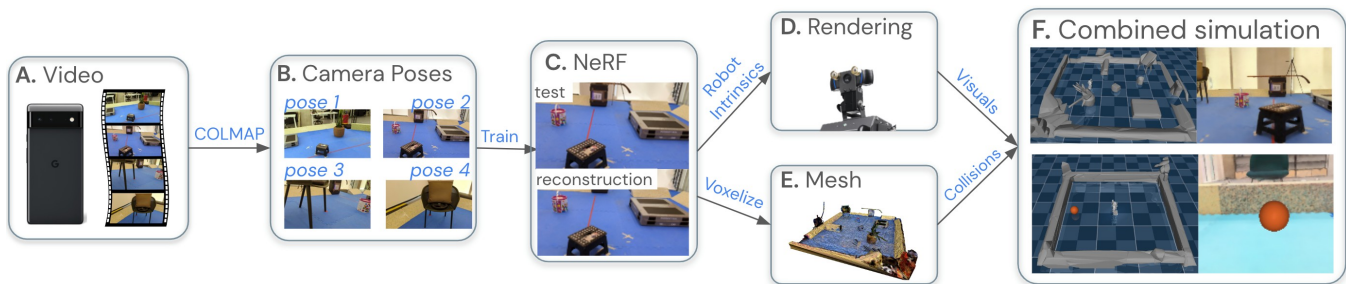


Fig. 2: Overview of our system for recreating a scene in a simulator. **A.** We collect a video of the scene using a generic phone. **B.** We use structure-from-motion software to label a subset of the video with camera poses. **C.** We train a NeRF on labeled images. **D.** We render the scene from novel views using the calibrated intrinsics of the robot’s head-mounted camera. **E.** We use the same NeRF to extract the scene geometry as a mesh. We coarsen the mesh and replace the floor with a flat primitive. **F.** We combine the simplified mesh with a model of a robot, and any other dynamic objects, in a physics simulator. See Fig. 3 for further details on this step.

B. Localization

Next, we extract $N \approx 1000$ keyframes from the captured video. We divide the video into N equal partitions, and for each partition we use the frame’s variance of the Laplacian to pick the least motion blurred frame. We use COLMAP [53]–[55], an open-source structure-from-motion (SfM) package, to estimate the intrinsics and extrinsics for each keyframe.

C. NeRF training

Given a dataset of images and corresponding camera poses, we train a NeRF [9] to render the scene from novel viewpoints. We use recent NeRF extensions for better reconstructions, improved reconstructed geometry, and decreased rendering times. To avoid artifacts while rendering at low resolutions, we sample the average of the volume over a normal distribution [56]. We use a space squashing formulation to support large capture areas, as well as a separate proposal network, and a distortion loss that encourages compact representations [10]. To improve the reconstructed geometry we optimise a separate specular and diffuse color [17].

NeRF rendering can be compute intensive even at low resolutions. While this is not critical in our context as we use the NeRF only for offline learning in simulation, to allow for faster experiment turnaround we implement a multi scale spatial hash grid approach [11]. This provides a significant speedup (order of magnitude), enabling rendering one frame in 6ms on a V100 GPU. We use a similar architecture as described in [11], adding a layer normalization [57] before the final MLP layer, and use swish activations [58] rather than ReLU activations. Additionally, we adapted the spatial hash grid approach [11] to allow sampling the radiance volume over a distribution. We blur training samples with a Gaussian blur with a random variance $\sigma_{blur} \in [\sigma_{min}, \sigma_{max}]$, and provide $\Sigma = \Sigma_{sample} * (1 + (\sigma_{blur} - \sigma_{min}))$ as an extra input to the final MLP [59]. This augmentation allows the network to interpolate samples in scale-space and improves our reconstruction significantly at lower resolutions (≈ 31.5 vs ≈ 35.4 average PSNR in a few held out images).

D. Rendering in simulation

The trained NeRF can be used for rendering the scene from novel viewpoints and camera intrinsics. In particular, we will use the intrinsic parameters of the robot’s camera

(Logitech C920) for rendering during agent training. We match camera intrinsics between sim and real by calibrating the robot’s camera, and use the obtained focal length and distortion parameters to render from the NeRF.

E. Collision mesh extraction

The NeRF learns a function to predict the radiance and density in space, i.e. the underlying scene geometry. We voxelize the predicted density, manually choose a threshold to convert it to a binary occupancy which is used to compute a mesh via the marching cubes algorithm [60]; this mesh is used only for collisions within our simulation.

The camera poses obtained from COLMAP, and hence also the collision mesh vertices, are expressed in an arbitrary reference frame (including an arbitrary scale). Therefore, we need to estimate a rigid transformation and scale between this frame of reference and the simulator’s world frame. We do this by solving a least-squares optimization that constrains the normal vector to the dominant floor plane in the mesh to be aligned with the simulator’s z-axis. We use Blender [61] to manually select points on the mesh’s floor for this purpose. We then manually rotate the mesh around the z-axis to a desired alignment with the simulator’s world frame and compute the relative scale between the NeRF and the world by comparing the size of an object within the mesh and the real world. We also replace the floor vertices in the mesh (which can have artifacts due to a lack of texture) with a flat plane. Lastly, for faster collision computation, we crop the mesh to the extents needed for simulation.

F. Physics simulation

We use MuJoCo [62] as our physics simulator. The simulated scene consists of the mesh extracted from NeRF attached to the world frame as a geometric primitive, the robot body, and possibly other dynamic objects (e.g. a ball).

Realistic rendering of such composite scenes is an active area of research [63]. We opt for a straightforward approach which is suitable for our tasks. We assume that the dynamic objects (all rendered with the MuJoCo built-in renderer) are always in the foreground of the static scene (rendered with NeRF); note that this doesn’t handle occlusions. Under this assumption the combined rendering is obtained by overlaying the dynamic objects rendering on top of the NeRF rendering (see Fig. 3 for a visualization).




Fig. 3: Our MuJoCo simulation is created by combining: (1) the learnt static scene mesh (Section III-E), (2) the dynamic object meshes and (3) the learnt static scene NeRF rendering (Section III-D) on which (4) the Mujoco rendering of dynamic objects (a ball and robot’s left arm in the camera image above) are overlaid. Other dynamic parameters (e.g. friction) are assumed known or measured.

IV. SIM2REAL TRAINING SETUP WITH NERF & MUJOCO

Once the combined MuJoCo simulation is set up we can train a policy purely in this simulation and deploy it directly in the real-world. We describe our training setup below.

A. Humanoid robot platform

We use a Robotis OP3 [64] robot for all our experiments. This low cost platform is a small humanoid (about 35 cm tall) with 20 actuated degrees of freedom (see Fig. 1). Actuators, and hence our learned policies, are operated in a position control mode with both D and I gains set to 0. Our policies run at 40Hz and rely on the robot’s on-board computer (2-core Intel NUC i3) and on-board sensors only. These include joint encoders, gyroscope, accelerometer, and a Logitech C920 camera attached to the robot’s head which is actuated via two joints attached to its torso. The gyroscope and accelerometer data are filtered at 125Hz to obtain an estimate of the gravity direction in the robot’s body frame using a Madgwick filter [65]. To encourage smoother movements, we apply an exponential filter with strength 0.8 to the control signals before passing them to the actuators.

B. Reducing the dynamics sim2real gap

We ensured that the robot’s sensors and actuators are modeled accurately in simulation. Specifically, we ensured that the simulated gyroscope and accelerometer data are low-pass filtered in the same way in simulation as on the real IMU chip. With these accurate models we found that policies trained in simulation transferred well to the real robot; hence we used only limited domain randomization. Particularly, we applied random pushes to the robot during training. We also applied constant delays per episode, sampled uniformly in the range of 10ms - 50ms as well as a 5 ms jitter to all simulated sensor data to reflect various latencies on the robot. At the beginning of each episode, we attach a random mass (up to 0.5kg) to a random position on the robot’s torso and randomize the IMU’s position on the torso. In tasks with a ball, we additionally randomize the ball’s mass (0.5 - 0.9kg) and radius (11.5 - 12.5cm) at the start of each episode (the real ball weighs 0.651kg with a radius of 12cm).

C. Regularizing policy learning for better sim2real transfer

Carefully choosing rewards for regularizing the robot’s behavior is important for successful transfer. In all of our tasks, we use the following reward components as a regularization:

1. a constant penalty whenever the robot’s yaw angular speed is larger than $\pi \text{ rad s}^{-1}$ to encourage slow turning;
2. L2 regularization on joint angles towards a default standing pose; and,
3. a walking reward encouraging the average of feet velocities in the robot’s forward direction to be 0.3 m s^{-1} . These rewards encourage the agent to learn gaits that transfer better, and also encourage better exploration.

D. Tasks

To demonstrate that our approach can scale to realistic scenes with complex geometries and support simple object interactions we choose two tasks:

Navigation and obstacle avoidance: We demonstrate our approach on a point to point visual navigation task where the robot has to reach multiple goals (specified as (x,y) coordinates in the NeRF’s frame of reference) while avoiding different obstacles such as a large plant, a chair, and walls; see Fig. 1 (left) and Fig. 5 (bottom) for a visualization of our scene which measures 5m x 4m. We chose three targets in different parts of the space that the robot has to reach. We automatically compute the free areas of the scene using the NeRF’s mesh and, during simulation, we randomly initialize the robot to a position and orientation within these areas.

The *reward* for training consists of the regularization terms described in Sec. IV-E, and two task-specific terms: 1. a sparse bonus upon reaching the goal location; 2. a walking reward like the one we use as a regularization but instead encouraging moving in the direction of the goal at a speed of 0.3 m s^{-1} . Episodes terminate whenever the robot’s body parts other than the feet touch the scene’s mesh. We consider an episode to be successful if the robot gets to $\leq 25 \text{ cm}$ of the target without falling and does not collide with any obstacles.

Ball pushing: As a proof of concept that we can combine static NeRF scenes with dynamic interacting objects, we consider a task in which the robot has to move a basketball to a corner of a 3m x 3m workspace (see Fig. 1, right). We model the basketball as a simple orange ball ignoring the fine black print which is barely visible at the resolutions we use (see Fig. 3 for an example simulated image). During training in simulation, each episode starts with the ball and robot randomly positioned. In half of all episodes, we initialize the ball just in front of the robot to speed up learning.

We again use the regularization terms in Sec. IV-E and two task-specific terms: 1. a reward for minimizing the distance between the ball and the goal region; and, 2. a reward for minimizing the distance between the robot and the ball if the ball is not moving towards the goal. Episodes are terminated whenever the robot falls. We consider an episode as successful if the robot gets the ball to the correct 1m x 1m corner square within 60 seconds. This task is much more challenging than navigation due to significant partial observability and interactions; the robot has to search for the ball, localize itself and the ball, and move it to the goal.

E. Policy training

All our policies are trained using DMPO [66], a state-of-the-art algorithm which combines distributional deep RL [67]




Fig. 4: The policy’s network architecture.

and MPO [68], [69]. Our policies take vision and proprioception as input; the policy network (Fig. 4) consists of a recurrent image encoder (to handle partial observability) which passes RGB camera images (30x40 or 60x80 resolution) through a small ResNet followed by an LSTM. The encoded images are then combined with a history of past 5 proprioceptive observations (gyroscope, accelerometer, gravity direction estimate, joint positions, and previous control signal) and passed through an MLP which outputs a diagonal Gaussian for sampling actions. In the visual navigation task we additionally include a one-hot encoding of the goal.

We use an asymmetric actor-critic setup for training in simulation where the critic, a separate neural network that is not evaluated on the robot, receives privileged information. Specifically, the critic shares the same network structure as the actor but we replace the image encoder with the simulation’s ground truth state (robot/object poses and velocities). This step is crucial for efficient learning in simulation.

Image augmentations: While the NeRF significantly reduces the sim2real gap with realistic scene renderings, we cannot easily modulate image intensity properties such as brightness or gain (see Fig. 1 (left) for comparison of real vs rendered images). Thus, we apply image augmentations during training: we randomize the brightness, saturation, hue, and contrast, and apply random translations to the image.

V. EXPERIMENTAL RESULTS

We now present our results. We highly encourage the reader to watch the accompanying video at <https://sites.google.com/view/nerf2real>.

A. Training time

We train policies in simulation and transfer them zero-shot to the robot for evaluation. Given ~ 1000 frames from a 5-6 minute video, COLMAP localization takes about 3-4 hours. Training the NeRF on this dataset takes about 20 minutes on a cluster of 8 V100s, though this could be significantly sped up further [11]. Mesh extraction and post-processing for setting up the simulation takes a few hours, and finally training the policies takes around 24 hours (about 8M gradient updates, and 128M environment steps) for navigation and twice as long for ball pushing.

B. Navigation and obstacle avoidance results

Evaluation setup & metrics: We compare the performance of learned goal-conditioned policies in simulation to zero-shot transfer performance in the real world. We use the following evaluation protocol: for each of the 3 goals we chose 3 unique initial positions and 3 orientations,

forming a total of 27 combinations. We perform two trials for each combination, for a total of 54 real-world episodes. We consider an episode to be successful if the robot reaches the goal (≤ 25 cm distance) without falling. We report the **overall success percentage**, i.e. the fraction of episodes the policy was successful, and the median **time taken to reach the target**. As the evaluation space is equipped with a motion capture system we can use this to compute the ground-truth position of the robot—this ground-truth information is only used for analysis and evaluation, not as an input to the policy.

As an analysis tool, our policies are trained together with an **auxiliary prediction network**, an MLP which predicts the robot’s belief about its 2D position and yaw based on the features output from the policy’s LSTM (we do not propagate any gradients from this loss). The belief is modelled as a mixture of five Gaussians and we use it to help disambiguate between policy failures due to confusing visual inputs and those due to other factors (e.g. the robot tripping). We use the difference between the mean of the belief distribution and the ground-truth position when quantifying the policy’s localization error which is averaged over the entire episode.

Results: Table I presents the results of our evaluation. We evaluated policies with two different input resolutions, 30x40 and 60x80; due to hardware limitations on the robot (2 CPU cores, no GPU) and the need to run a policy step within 25ms, we were unable to evaluate higher resolutions.

We draw attention to a few interesting trends: 1) **Our policies do not exhibit a significant gap between performance in simulation and on the real robot.** On the real robot, the 60x80 policy successfully reached the goal in **47/54 episodes (87 \pm 5%)**. The lower resolution 30x40 policy performs slightly worse and was successful in 37/54 episodes (69 \pm 6%). This is remarkably similar to the performance of these policies in simulation. Based on monitoring the policy’s belief about its pose, we estimate that about half of the failures of the 40x30 policy were due to collisions with obstacles and/or the robot falling down, and the rest due to localization failures, but it is impossible to perfectly disambiguate failure modes. 2) The policies take **similar amounts of time to reach the goal in simulation and on the robot** demonstrating that dynamical properties of the behavior such as the gait velocity also do not suffer from a sim2real gap. 3) As expected, our policies learn representations which are informative about the robot’s current pose; these transfer well to the real-world. Using the **auxiliary prediction network**, we visualized the agent’s belief and evaluated localization errors in simulation and the real world. The median localization error across all control steps and trials is ~ 0.25 m. We saw that **several failures of the policy correlated with higher localization error**, specifically on a single target near the potted plant. This was particularly evident for the 40x30 policy (0.33m for failures vs 0.23m for successes). We show a visualization of the belief for a successful trajectory in Fig. 5, which demonstrates that the agent is able to quickly localize itself accurately with ~ 2 seconds of data and rarely loses track. Videos of simulated and real-world evaluations can be found on our [website](#).




Fig. 5: Localization performance of the agent. Top row: Robot camera images from a successful zero-shot transfer trial. The target is next to the potted plant. Bottom row: Visualization of the robot’s belief over it’s 2D position in the scene, shown as a heatmap on a top-down view of the scene. **White X**: Ground truth position from motion capture (not input to policy); **Green Triangle**: Target position.

Policy resolution	Simulation results			Zero-shot transfer results		
	Success	Time taken	Localization error	Success	Time taken	Localization error
30 x 40	73±3%	10.3±0.3 sec	0.17±0.01 m	69±6%	10.3±0.4 sec	0.24±0.01 m
60 x 80	86±2%	10.8±0.1 sec	0.19±0.004 m	87±5%	11.2±0.4 sec	0.27±0.04 m

TABLE I: Sim and real performance (with standard errors) of the trained policies on the navigation and obstacle avoidance task. Time taken and Localization error values are median statistics across all evaluation episodes.

Policy resolution	Sim Success		Real Success	
	Center	Wall	Center	Wall
30 x 40	99%	100%	78±7%	43±14%

TABLE II: Sim and real performance (with standard errors) of trained policies on the ball pushing task.

C. Ball pushing results

Evaluation setup & metrics: Similar to the navigation task, we compare the **average success** of policies between simulation and real-world. We consider two different evaluation setups in the real world, a **Center** setting where the ball is initialized in the center cell of the workspace (Fig. 1, right) and the robot is initialized in the center of any of the eight cells near the wall with 4 different orientations per cell (32 episodes total), and a **Wall** setting where the robot is initialized in the center cell facing the target corner, and the ball is initialized near the center of one of the remaining 7 cells (except the target cell) with two trials each (14 episodes total). The target is always fixed to be the corner with the two cones (see Fig. 1, right), and an episode is considered successful if the robot can get the ball to the 1m x 1m target cell within 60 seconds without falling down.

Results: Table II presents the evaluation results. We highlight two key points: 1) As can be seen in the accompanying video, our policies use the robot’s hands to move the ball, and exhibit active perception when searching and tracking the ball. These behaviors emerge from the task requirements and are **not explicitly encouraged by the reward**; they also transfer successfully to the real world. 2) While our policies perform well, the sim2real gap is larger for this contact-rich task. This is especially true for the **Wall** initializations; unless the robot executes a perfect push, the ball often gets stuck near the walls and the robot struggles to move it.

Videos of all our results, and comparisons of renderings from the NeRF, COLMAP reconstructions (which we show as a baseline) and real images can also be found on our [website](#); renderings from COLMAP and textures on top of the NeRF mesh show artifacts, especially on the background, which will significantly inhibit sim2real transfer.

VI. DISCUSSION AND LIMITATIONS

We presented a pipeline for creating simulation environments of visually complex scenes in a way that allows training vision guided policies for sim2real transfer. To this end we combine the scene geometry and rendering function derived from a NeRF with a known physics model of the robot and (optionally) additional objects. In principle, our approach is embodiment independent and it can be automated further in future work. For instance, new NeRF-like models [17], [70], [71] may improve scene geometry reconstruction thus eliminating the need for manual postprocessing of minimally textured areas like the floor. Evaluating the approach on contact-rich tasks such as climbing on objects will allow us to better assess the current limitations of the approach, and may guide future NeRF and scene modeling developments.

We have currently opted for a very simple approach to composing the rendering of an a priori known object with the rendering of a static NeRF scene. However, the field has been actively working on NeRF-like approaches to photorealistic rendering of composite scenes [72], including ways for segmenting a static scene into dynamic objects and representing each using a separate NeRF [16]. If necessary, our pipeline can leverage any of these improvements (and will have to for more complex dynamic scenes). Similarly, we believe that recent progress on eliminating the computationally expensive localization step from NeRF pipelines [18], [19], and speeding up both NeRF [11] and RL training [73] will soon enable going from a video of a scene to a trained policy within minutes or hours instead of 1-2 days, potentially enabling running our setup online on the robot during deployment.

Impact statement: This work presents an approach to train vision guided policies for general robotics systems. While in its current form the approach is unlikely to enable real world applications, future research may make possible a range of applications that can benefit humanity. We strongly oppose any applications designed to bring harm to humans.

Acknowledgements: We thank Neil Sreendra, Marlon Gwira, Kushal Patel, Nathan Batchelor, Federico Casarini, Jon Scholz, Francesco Romano and Claudio Fantacci.

REFERENCES

- [1] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Sci Robot*, vol. 4, Jan. 2019.
- [2] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Science Robotics*, vol. 7, no. 62, p. eabk2822, 2022.
- [3] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, *et al.*, "Solving rubik's cube with a robot hand," *arXiv preprint arXiv:1910.07113*, 2019.
- [4] P. Anderson, A. Shrivastava, J. Truong, A. Majumdar, D. Parikh, D. Batra, and S. Lee, "Sim-to-real transfer for vision-and-language navigation," in *Conference on Robot Learning*, pp. 671–681, PMLR, 2021.
- [5] 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.
- [6] 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*, pp. 9068–9079, 2018.
- [7] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijnmans, 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.
- [8] A. Szot, A. Clegg, E. Undersander, E. Wijnmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets, *et al.*, "Habitat 2.0: Training home assistants to rearrange their habitat," *Advances in Neural Information Processing Systems*, vol. 34, pp. 251–266, 2021.
- [9] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *CoRR*, vol. abs/2003.08934, 2020.
- [10] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, "Mip-nerf 360: Unbounded anti-aliased neural radiance fields," *CoRR*, vol. abs/2111.12077, 2021.
- [11] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Trans. Graph.*, vol. 41, pp. 102:1–102:15, July 2022.
- [12] C. Reiser, S. Peng, Y. Liao, and A. Geiger, "Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps," *CoRR*, vol. abs/2103.13744, 2021.
- [13] M. Tancik, V. Casser, X. Yan, S. Pradhan, B. Mildenhall, P. P. Srinivasan, J. T. Barron, and H. Kretschmar, "Block-nerf: Scalable large scene neural view synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8248–8258, 2022.
- [14] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer, "D-nerf: Neural radiance fields for dynamic scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10318–10327, 2021.
- [15] K. Park, U. Sinha, J. T. Barron, S. Bouaziz, D. B. Goldman, S. M. Seitz, and R. Martin-Brualla, "Nerfies: Deformable neural radiance fields," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5865–5874, 2021.
- [16] K. Stelzner, K. Kersting, and A. R. Kosiorek, "Decomposing 3d scenes into objects via unsupervised volume segmentation," *arXiv preprint arXiv:2104.01148*, 2021.
- [17] D. Verbin, P. Hedman, B. Mildenhall, T. E. Zickler, J. T. Barron, and P. P. Srinivasan, "Ref-nerf: Structured view-dependent appearance for neural radiance fields," *CoRR*, vol. abs/2112.03907, 2021.
- [18] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu, "Nerf: Neural radiance fields without known camera parameters," *arXiv preprint arXiv:2102.07064*, 2021.
- [19] E. Sucar, S. Liu, J. Ortiz, and A. J. Davison, "imap: Implicit mapping and positioning in real-time," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6229–6238, 2021.
- [20] Z. Chen, T. Funkhouser, P. Hedman, and A. Tagliasacchi, "Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures," *arXiv preprint arXiv:2208.00277*, 2022.
- [21] L. Yen-Chen, P. Florence, J. T. Barron, A. Rodriguez, P. Isola, and T.-Y. Lin, "inert: Inverting neural radiance fields for pose estimation," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1323–1330, IEEE, 2021.
- [22] L. Yen-Chen, P. Florence, J. T. Barron, T.-Y. Lin, A. Rodriguez, and P. Isola, "Nerf-supervision: Learning dense object descriptors from neural radiance fields," *arXiv preprint arXiv:2203.01913*, 2022.
- [23] J. Ichnowski, Y. Avigal, J. Kerr, and K. Goldberg, "Dex-nerf: Using a neural radiance field to grasp transparent objects," *arXiv preprint arXiv:2110.14217*, 2021.
- [24] D. Driess, Z. Huang, Y. Li, R. Tedrake, and M. Toussaint, "Learning multi-object dynamics with compositional neural radiance fields," *arXiv preprint arXiv:2202.11855*, 2022.
- [25] M. Adamkiewicz, T. Chen, A. Caccavale, R. Gardner, P. Culbertson, J. Bohg, and M. Schwager, "Vision-only robot navigation in a neural radiance world," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4606–4613, 2022.
- [26] S. K. Ramakrishnan, A. Gokaslan, E. Wijnmans, O. Maksymets, A. Clegg, J. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, *et al.*, "Habitat-matterport 3d dataset (hm3d): 1000 large-scale 3d environments for embodied ai," *arXiv preprint arXiv:2109.08238*, 2021.
- [27] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijnmans, 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*, pp. 9339–9347, 2019.
- [28] B. Shen, F. Xia, C. Li, R. Martín-Martín, L. Fan, G. Wang, C. Pérez-D'Arpino, S. Buch, S. Srivastava, L. Tchammi, *et al.*, "igibson 1.0: a simulation environment for interactive tasks in large realistic scenes," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7520–7527, IEEE, 2021.
- [29] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, D. Gordon, Y. Zhu, A. Gupta, and A. Farhadi, "Ai2-thor: An interactive 3d environment for visual ai," *arXiv preprint arXiv:1712.05474*, 2017.
- [30] M. Deitke, W. Han, A. Herrasti, A. Kembhavi, E. Kolve, R. Mottaghi, J. Salvador, D. Schwenk, E. VanderBilt, M. Wallingford, *et al.*, "Robothon: An open simulation-to-real embodied ai platform," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3164–3174, 2020.
- [31] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, *et al.*, "On evaluation of embodied navigation agents," *arXiv preprint arXiv:1807.06757*, 2018.
- [32] E. Wijnmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, "Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames," *arXiv preprint arXiv:1911.00357*, 2019.
- [33] T. Chen, S. Gupta, and A. Gupta, "Learning exploration policies for navigation," *arXiv preprint arXiv:1903.01959*, 2019.
- [34] A. Khandelwal, L. Weihs, R. Mottaghi, and A. Kembhavi, "Simple but effective: Clip embeddings for embodied ai," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14829–14838, 2022.
- [35] D. Batra, A. Gokaslan, A. Kembhavi, O. Maksymets, R. Mottaghi, M. Savva, A. Toshev, and E. Wijnmans, "Objectnav revisited: On evaluation of embodied agents navigating to objects," *arXiv preprint arXiv:2006.13171*, 2020.
- [36] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. Van Den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3674–3683, 2018.
- [37] J. Truong, S. Chernova, and D. Batra, "Bi-directional domain adaptation for sim2real transfer of embodied navigation agents," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2634–2641, 2021.
- [38] J. Truong, D. Yarats, T. Li, F. Meier, S. Chernova, D. Batra, and A. Rai, "Learning navigation skills for legged robots with learned robot embeddings," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 484–491, IEEE, 2021.
- [39] Z. Fu, A. Kumar, A. Agarwal, H. Qi, J. Malik, and D. Pathak, "Coupling vision and proprioception for navigation of legged robots," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17273–17283, 2022.
- [40] 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 *2017 IEEE/RSJ international con-*

- ference on intelligent robots and systems (IROS), pp. 23–30, IEEE, 2017.
- [41] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, “Unsupervised pixel-level domain adaptation with generative adversarial networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3722–3731, 2017.
- [42] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, “Closing the sim-to-real loop: Adapting simulation randomization with real world experience,” in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8973–8979, IEEE, 2019.
- [43] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, et al., “Using simulation and domain adaptation to improve efficiency of deep robotic grasping,” in *2018 IEEE international conference on robotics and automation (ICRA)*, pp. 4243–4250, IEEE, 2018.
- [44] A. X. Lee, C. M. Devin, Y. Zhou, T. Lampe, K. Bousmalis, J. T. Springenberg, A. Byravan, A. Abdolmaleki, N. Gileadi, D. Khosid, et al., “Beyond pick-and-place: Tackling robotic stacking of diverse shapes,” in *5th Annual Conference on Robot Learning*, 2021.
- [45] F. Sadeghi and S. Levine, “Cad2rl: Real single-image flight without a single real image,” *arXiv preprint arXiv:1611.04201*, 2016.
- [46] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, 2019.
- [47] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science Robotics*, vol. 5, p. eabc5986, Oct 2020.
- [48] X. B. Peng, E. Coumans, T. Zhang, T.-W. Lee, J. Tan, and S. Levine, “Learning agile robotic locomotion skills by imitating animals,” *arXiv preprint arXiv:2004.00784*, 2020.
- [49] S. Bohez, S. Tunyasuvunakool, P. Brakel, F. Sadeghi, L. Hasenclever, Y. Tassa, E. Parisotto, J. Humplik, T. Haarnoja, R. Hafner, M. Wulfmeier, M. Neunert, B. Moran, N. Siegel, A. Huber, F. Romano, N. Batchelor, F. Casarini, J. Merel, R. Hadsell, and N. Heess, “Imitate and repurpose: Learning reusable robot movement skills from human and animal behaviors,” *arXiv*, Mar. 2022.
- [50] W. Yu, V. C. Kumar, G. Turk, and C. K. Liu, “Sim-to-real transfer for biped locomotion,” *arXiv preprint arXiv:1903.01390*, 2019.
- [51] Z. Li, X. Cheng, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, “Reinforcement learning for robust parameterized locomotion control of bipedal robots,” *arXiv preprint arXiv:2103.14295*, 2021.
- [52] J. Siekmann, K. Green, J. Warila, A. Fern, and J. Hurst, “Blind bipedal stair traversal via sim-to-real reinforcement learning,” *arXiv preprint arXiv:2105.08328*, 2021.
- [53] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [54] J. L. Schönberger, E. Zheng, M. Pollefeys, and J.-M. Frahm, “Pixel-wise view selection for unstructured multi-view stereo,” in *European Conference on Computer Vision (ECCV)*, 2016.
- [55] J. L. Schönberger, T. Price, T. Sattler, J.-M. Frahm, and M. Pollefeys, “A vote-and-verify strategy for fast spatial verification in image retrieval,” in *Asian Conference on Computer Vision (ACCV)*, 2016.
- [56] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan, “Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields,” *CoRR*, vol. abs/2103.13415, 2021.
- [57] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” 2016.
- [58] P. Ramachandran, B. Zoph, and Q. V. Le, “Searching for activation functions,” *CoRR*, vol. abs/1710.05941, 2017.
- [59] H. Baatz, J. Granskog, M. Papas, F. Rousselle, and J. Novák, “Nerf-text: Neural reflectance field textures,” in *Eurographics Symposium on Rendering*, The Eurographics Association, June 2021.
- [60] W. E. Lorensen and H. E. Cline, “Marching cubes: A high resolution 3d surface construction algorithm,” *ACM siggraph computer graphics*, vol. 21, no. 4, pp. 163–169, 1987.
- [61] B. O. Community, *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018.
- [62] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, IEEE, 2012.
- [63] B. Yang, Y. Zhang, Y. Xu, Y. Li, H. Zhou, H. Bao, G. Zhang, and Z. Cui, “Learning object-compositional neural radiance field for editable scene rendering,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13779–13788, 2021.
- [64] “Robotis op3.” <https://emanual.robotis.com/docs/en/platform/op3/introduction/>.
- [65] S. Madgwick et al., “An efficient orientation filter for inertial and inertial/magnetic sensor arrays,” *Report x-io and University of Bristol (UK)*, vol. 25, pp. 113–118, 2010.
- [66] M. Hoffman, B. Shahriari, J. Aslanides, G. Barth-Maron, F. Behbahani, T. Norman, A. Abdolmaleki, A. Cassirer, F. Yang, K. Baumli, S. Henderson, A. Novikov, S. G. Colmenarejo, S. Cabi, C. Gulcehre, T. L. Paine, A. Cowie, Z. Wang, B. Piot, and N. de Freitas, “Acme: A research framework for distributed reinforcement learning,” *arXiv preprint arXiv:2006.00979*, 2020.
- [67] M. G. Bellemare, W. Dabney, and R. Munos, “A distributional perspective on reinforcement learning,” in *International Conference on Machine Learning*, pp. 449–458, PMLR, 2017.
- [68] A. Abdolmaleki, J. T. Springenberg, Y. Tassa, R. Munos, N. Heess, and M. Riedmiller, “Maximum a posteriori policy optimisation,” 2018.
- [69] A. Abdolmaleki, J. T. Springenberg, J. Degraeve, S. Bohez, Y. Tassa, D. Belov, N. Heess, and M. Riedmiller, “Relative entropy regularized policy iteration,” 2018.
- [70] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman, “Volume rendering of neural implicit surfaces,” *CoRR*, vol. abs/2106.12052, 2021.
- [71] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, “Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction,” *CoRR*, vol. abs/2106.10689, 2021.
- [72] M. Guo, A. Fathi, J. Wu, and T. Funkhouser, “Object-centric neural scene rendering,” *arXiv preprint arXiv:2012.08503*, 2020.
- [73] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to walk in minutes using massively parallel deep reinforcement learning,” in *5th Annual Conference on Robot Learning*, 2021.