

Success Weighted by Completion Time: A Dynamics-Aware Evaluation Criteria for Embodied Navigation

Naoki Yokoyama¹ Sehoon Ha^{1,2} Dhruv Batra^{1,3}

Abstract—We present Success weighted by Completion Time (SCT), a new metric for evaluating navigation performance for mobile robots. Several related works on navigation have used Success weighted by Path Length (SPL) as the primary method of evaluating the path an agent makes to a goal location, but SPL is limited in its ability to properly evaluate agents with complex dynamics. In contrast, SCT explicitly takes the agent’s dynamics model into consideration, and aims to accurately capture how well the agent has approximated the fastest navigation behavior afforded by its dynamics. While several embodied navigation works use point-turn dynamics, we focus on unicycle-cart dynamics for our agent, which better exemplifies the dynamics model of popular mobile robotics platforms (e.g., LoCoBot, TurtleBot, Fetch, etc.). We also present RRT*-Unicycle, an algorithm for unicycle dynamics that estimates the fastest collision-free path and completion time from a starting pose to a goal location in an environment containing obstacles. We experiment with deep reinforcement learning and reward shaping to train and compare the navigation performance of agents with different dynamics models. In evaluating these agents, we show that in contrast to SPL, SCT is able to capture the advantages in navigation speed a unicycle model has over a simpler point-turn model of dynamics. Lastly, we show that we can successfully deploy our trained models and algorithms outside of simulation in the real world. We embody our agents in a real robot to navigate an apartment, and show that they can generalize in a zero-shot manner. A video summary is available here: <https://youtu.be/QOQ56XVIYVE>

I. INTRODUCTION

Mobile robots have the potential to automate several tasks in the real world and alleviate the amount of work to be done by humans. But for mobile robots to be effective in the physical world, they must be able to effectively navigate previously unseen environments. Yet, without access to a map or floorplan, inferring an efficient path to a distant goal in a novel environment using only the robot’s onboard sensors is a challenging task.

Several recent works have shown that deep reinforcement learning (DRL) can be used to teach virtual robots in simulators to accomplish tasks such as navigating to a goal point ([4], [9], [17], [20]), or searching for a distant object ([2], [3]), even within previously unseen environments. However, these virtual robots are typically modeled as idealized cylinders, equipped with a limited high-level action space (move-forward 0.25m, pivot $\pm 10^\circ$). A subtle but important

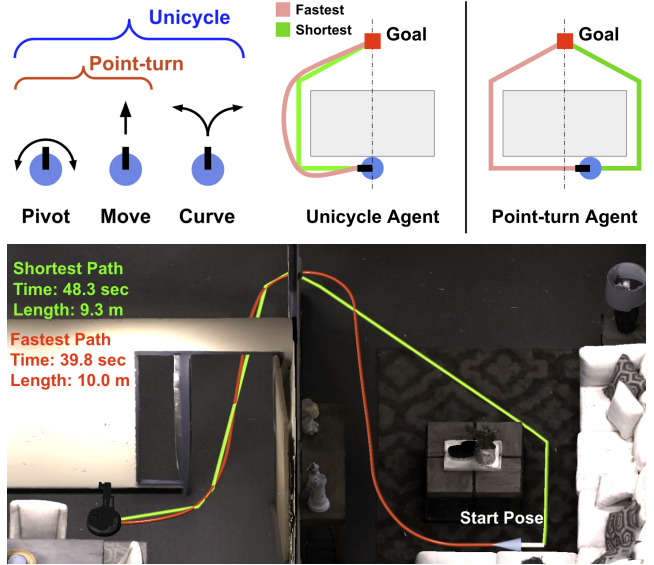


Fig. 1: The shortest path is not always the fastest. *Top left*: In addition to pivoting and moving straight like a point-turn model, a unicycle-cart dynamics model can move forward and turn (curve) simultaneously. *Top right*: A path faster than the shortest path may exist for both types of dynamics. The fastest path depends on the agent’s maximum linear and angular velocities, and initial heading. *Bottom*: Fastest path (calculated with RRT*-Unicycle) and shortest path visualized in Habitat, where our agents are trained.

consequence of the forced-choice nature of such an action-space is that these idealized agents are assumed to have ‘point-turn’ dynamics (illustrated in Figure 1), in which they must come to a full halt to turn. Deployment of such policies to real mobile platforms such as in [9] and [12] results in motion that is often jerky and time-consuming, and can result in increased power consumption or localization errors that compound from abrupt motions, as shown in [13].

To address these challenges and navigate more efficiently, we propose to use a more complex model of dynamics, namely, a unicycle-cart model. An agent with unicycle dynamics can not only pivot in-place and move straight like a point-turn agent, but can also make smooth curves by moving with both linear and angular velocity. This allows it to execute smoother turns without halting to navigate around obstacles more quickly than a point-turn agent. These points are illustrated in Figure 1.

However, we find that navigation performance metrics such as Success weighted by Path Length (SPL) are unable to properly capture the benefits in speed brought by more complex dynamics models, and instead favors the jerky paths produced by point-turn behavior. Therefore, we propose a

¹NY, SH, DB are with the College of Electrical Engineering and College of Computing, Georgia Institute of Technology, Atlanta, GA 30318 USA. ²SH is with Robotics at Google. ³DB is with Facebook AI Research (FAIR). The Georgia Tech effort was supported in part by NSF, DARPA, ONR YIP, and ARO PECASE. (e-mail: nyokoyama@gatech.edu; sehoonha@gatech.edu; dbatra@gatech.edu)

new metric, Success weighted by Completion Time (SCT), which addresses the shortcomings of navigation metrics such as SPL by explicitly taking into consideration the dynamics model of the agent. To calculate SCT, we formulate a method called RRT*-Unicycle to find the fastest path from a start to a goal for a unicycle model of dynamics, which can be extended to other models of dynamics.

In this work, we use DRL to train agents with unicycle-cart dynamics to successfully navigate to the goal location significantly faster than those trained with point-turn dynamics, despite having the same maximum linear and angular velocities (i.e., not being able to move any faster). We show that SPL fails to capture the faster successful completion speeds of the unicycle agents, while SCT can. We also propose a decaying scheme for reward shaping, which we show results in agents learning faster navigation behavior than a fixed shaped reward function. Lastly, we deploy a unicycle agent on a real robot, and show that our model and reward scheme lead to trajectories that are 35% faster than those of the point-turn agent with identical maximum linear and angular velocities.

II. RELATED WORK

Visual Navigation with DRL. Several works have also explored the use of DRL in training visual navigation policies using photorealistic 3D simulators of indoor environments, such as [3], [6], [12], [16], [18], [20]. However, all of these works consider agents that choose from a small discrete set of actions, which only allow the agent to move forward or pivot in-place. Furthermore, many of these works evaluate their agents based on Success weighted by Path Length (SPL), a metric Anderson et. al recommended to adopt as the primary method of evaluating navigation performance in [1]. In this work, we show why SPL is not optimal for evaluating agents with more complex dynamics, and propose a new metric to evaluate the navigation performance for such agents.

Evaluation Metrics. The metric proposed by [5], Success weighted by Number of Actions, closely resembles our proposed metric. Like our metric, it also penalizes redundant idle actions that do not necessarily increase the length of the agent’s path (e.g., excessively turning in-place). However, it ultimately derives the baseline number of actions to compare against using the shortest path to the goal, while we focus on the fastest path conditioned on the agent’s dynamics.

III. PROBLEM FORMULATION

A. Task: Point Goal Navigation

For this work, we consider the task of PointGoal Navigation (PointNav), as detailed in [1]. In this task, the robot is initialized within a previously unseen environment, and must navigate to a goal point specified in relative coordinates. An *episode* is defined by the environment to navigate, a starting location and heading, and a goal location. The agent must reach the goal point using only its onboard sensors as input. Our agents do not have access to a map; all maps and calculated paths shown in this work are for performance evaluation and illustrative purposes only.

Success criteria. Episodes are considered successful only if the agent is both within a certain radius of the goal point and invokes an action that terminates the episode. We choose the radius of LoCoBot [7] (0.2m) as our success radius because it is the mobile platform that we use in our real-world experiments. Episodes are terminated and unsuccessful after 500 action steps if the agent has not terminated it yet.

B. The limitations of SPL

SPL scores the length of an agent’s path based on how close it is to the length of the shortest path. Specifically,

$$\text{SPL} = S \frac{L}{\max(P, L)} \quad (1)$$

where S is 0 or 1 depending on whether the agent successfully completed the episode, P is the length of the agent’s path, and L is the length of the shortest path from the start point to the goal point.

However, the shortest path to a goal is not necessarily the fastest. In many cases, the shortest path can actually be much slower or energy inefficient than an alternative path, as depicted in Figure 1. The fastest path depends highly on the agent’s dynamics, as well as its initial heading at the start of the episode. Because SPL only evaluates the agent by comparing the length of its path to the shortest path length from the start to the goal, it incentivizes tight turns around obstacles without consideration for how inefficient following such a path may be for the agent’s dynamics. This unfairly favors simpler point-turn behaviors of motion, and discourages agents to leverage their model of dynamics in any way that would deviate from that of a point-turn. Additionally, it does not penalize unnecessary or redundant actions that do not lengthen the agent’s path, such as pausing or pivoting. This and other limitations of SPL are discussed in Batra et al [2].

C. Success weighted by Completion Time

To overcome all of these limitations of SPL, we propose Success weighted by Completion Time (SCT), defined as:

$$\text{SCT} = S \frac{T}{\max(C, T)} \quad (2)$$

where C is the agent’s completion time, and T is the shortest possible amount of time it takes to reach the goal point from the start point while circumventing obstacles based on the agent’s dynamics.

One alternative to SCT we considered was simply computing the agent’s average completion time on successful episodes. However, computing the lower and upper bounds of the average completion time is not trivial, making it difficult to tell how well the robot *could* have performed the task. It is also not comparable across datasets; for example, if one dataset contains navigation environments larger than another dataset’s on average, its average successful completion time would likely be longer in comparison. SCT addresses these issues by measuring how close an agent’s completion speed is to being optimal, using pre-computed fastest paths to compare the agent’s paths against.

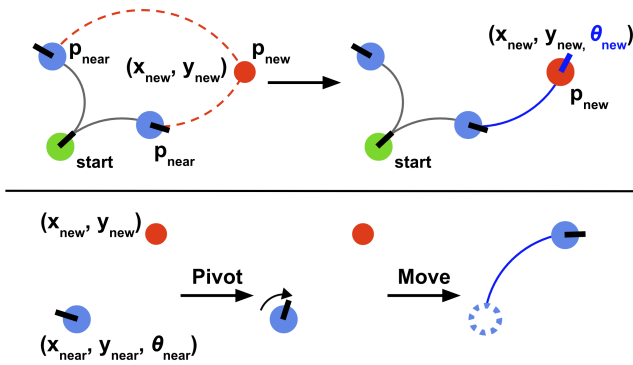


Fig. 2: *Top*: The heading (θ) of a new node is determined by the fastest path from the start that can be made from existing neighbors. *Bottom*: We calculate the fastest path from a pose to a target location in free space by initially pivoting the agent towards the target with max angular velocity (if optimal), before moving towards the target location with max linear velocity and constant angular velocity.

D. Calculating the fastest path time: RRT*-Unicycle

To compute SCT, we need to compute the fastest possible completion time of an episode. Although a navigation mesh, an auxiliary data structure often provided by simulators representing the navigable areas of an environment, greatly simplifies the process of finding shortest paths using A*, they do not encode anything about a given agent’s dynamics. Therefore, it cannot be easily exploited to find the fastest path. For such a task, creating a new graph using RRT* [10] that considers the agent’s dynamics is often the preferred approach [15], [19].

We develop an adapted version of RRT*, which we call RRT*-Unicycle, to find the fastest path from the start to the goal of an episode that is conditioned on the robot’s dynamics. Recall that RRT* is a sampling-based optimal motion-planning algorithm that, given a map, finds an obstacle-free path between two points that minimizes a cost, and is known to approach the optimal solution as the number of samples approaches infinity. The cost we aim to minimize is the time it takes the agent to travel from its starting pose to the goal point. RRT* works by creating a graph that traverses the free space of an environment, constructing and optimizing collision-free paths between sampled points until a full path from the start to the goal can be assembled. A full review of RRT* is beyond the scope of this document, and we refer the reader to [10]. For the purpose of describing RRT*-Unicycle, we will be highlighting the changes made to the sampling, connecting, and rewiring stages of RRT*.

Selecting a new node. RRT*-Unicycle operates in 2D space, where each node p in the constructed graph represents a pose of the agent (x, y, θ) , defined by a location and heading. When sampling a new node p_{new} , we initially leave θ_{new} uninitialized; only (x_{new}, y_{new}) is fixed, which is randomly sampled from the free space. However, if p_{new} is farther than a certain step distance r from any existing node, we project it on to the circle with radius r around the closest existing node. If p_{new} is no longer in the free space (i.e., within an obstacle), we discard it and sample again.

Similar to RRT*-Smart [14], once a viable path from the

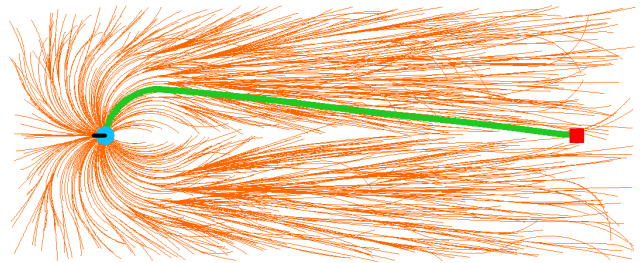


Fig. 3: Our method of finding the fastest path can link nearby nodes with circular arcs to approximate faster, more complex splines.

start to the goal is found, we bias our sampling to the free space near this path to more rapidly approach the optimal solution. Additionally, when a viable path has not been found yet, we bias sampling around points on the shortest path (pre-computed using A* on the navigation mesh) to more efficiently explore the free space until a viable path has been found. We avoid explicitly adding the shortest path to the graph, as it is often unlikely to be the fastest path for a unicycle model.

Connecting to an existing node. Once p_{new} is sampled, for each existing node p_{near} within a certain radius of p_{new} , we calculate the fastest path from $(x_{near}, y_{near}, \theta_{near})$ to (x_{new}, y_{new}) , as explained in the next subsection. The p_{near} that yields the fastest path from the start to p_{new} without colliding into any obstacles is chosen as the parent of p_{new} . To determine if a collision would occur, we uniformly sample points along a potential path and ensure that none intersect with an obstacle. When p_{new} is finally added to the graph, θ_{new} is set as the heading of the agent when it arrived at location (x_{new}, y_{new}) from the chosen p_{near} via the fastest calculated path, as shown in the upper portion of Figure 2.

Leaving θ_{new} initially unfixed simplifies calculating the fastest path from a p_{near} to p_{new} , as it is easier to calculate the fastest path from a source pose (x_s, y_s, θ_s) to a target location (x_t, y_t) , than to a target pose (x_t, y_t, θ_t) . We find that allowing θ_{new} to instead be set by the fastest path from the neighboring node offering the shortest cost from the start point converges much faster towards the optimal path.

Calculating the fastest path in free space. To calculate the fastest path from a pose to a nearby location (assuming there are no obstacles), for simplicity, we assume that the fastest behavior is to follow an arc with maximum linear velocity and a constant angular velocity that would lead the agent directly to the target location from its current pose. An agent can perform an in-place pivot with maximum angular velocity towards the target location to begin the path only if it would decrease the total path completion time. To determine how much pivoting yields the shortest path time, we used a brute-force method to generate a lookup table mapping the heading and distance relative to the target location to the optimal amount of pivoting needed before moving along the arc. This behavior is shown in the bottom of Figure 2. Though the resulting circular arc is not always the fastest path, especially over longer distances, it is a good close approximation for nearby points. By joining enough nearby

points together with this method, we can form the more complex splines between points in free space that represent faster paths than an arc, as shown in Figure 3.

RRT* Rewiring. During the rewiring step, the nodes near p_{new} disconnect from their parent node and adopt p_{new} as its new parent if doing so decreases their cost to be reached from start node. Thus, we must calculate the cost from p_{new} to a nearby p_{near} , where both nodes now have fixed θ values. In other words, the fastest path time must be calculated from a pose to a pose, rather than from a pose to a location. To calculate this cost, we simply calculate the fastest path and its completion time from the source pose to the target poses’s location, as described previously, and then add the amount of time necessary to pivot in-place until θ_{near} is reached as fast as possible. Though this is an overestimation of the true fastest path between two poses, we find that this method is still able to lead to a near-optimal path with enough samples.

Extracting the best path. To get the best path to the goal from the tree, for each node p_{near} within a certain radius of (x_{goal}, y_{goal}) , we compute the cumulative fastest path completion time from the start node p_{start} to p_{near} , add the completion time of the fastest path from p_{near} to (x_{goal}, y_{goal}) , and select the p_{near} that minimizes this sum. The path from p_{start} to the selected p_{near} joined with the fastest path from p_{near} to (x_{goal}, y_{goal}) represents the fastest path from the start pose to the goal location found by the tree.

Summary. The full procedure is summarized in Algorithm 1. An example of a constructed tree where the fastest paths to all sampled nodes are visualized is shown in Figure 4, where the fastest path to the goal point is highlighted. Notice that the fastest path is not the shortest.

Though we focus RRT*-Unicycle on a unicycle-cart model of dynamics, which reflects the type of motion widely used by real indoor mobile robots, it can easily be extended to other dynamics models as well. This can be done by changing the way in which the fastest path from a pose to a target location under an obstacle-free setting is calculated, which defines the cost of each edge in the graph.

IV. EXPERIMENTAL SETUP

A. Simulator and dataset

We train our agents within Habitat [17], a high-performance physics-enabled 3D simulator that provides the virtual robot with first-person photorealistic visual data. We use the same simulator settings as the Habitat 2019 challenge [17]. For training and evaluation, we use the Gibson-4+ dataset [17], comprised of 86 high quality 3D scans and meshes curated from the full Gibson dataset [21]. The Gibson dataset includes scans of various indoor environments, such as apartments, multi-level homes, offices, houses, hospitals, and gyms. The Gibson-4+ dataset only contains scenes that have been manually rated 4 or above on a quality scale from 0 to 5, and are free of significant reconstruction artifacts such as holes in texture or cracks in floor surfaces. In total, the training split of Gibson-4+ contains 3.6M episodes distributed across the 72 scenes, while the validation split

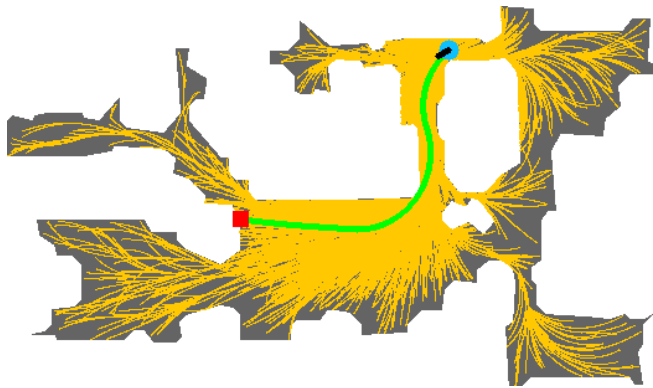


Fig. 4: A tree created with RRT*-Unicycle for calculating the fastest path. Fastest found paths to all sampled points shown in yellow, and fastest found path to the goal in green. These paths consider the maximum linear and angular velocities of our unicycle-cart agents, as well as its initial heading at the start of the episode. Neither the maps nor the paths are provided to the agents during training or testing; they are for performance evaluation purposes only.

Algorithm 1: RRT*-Unicycle

```

1  $E \leftarrow 0.2$  // Exploration probability
2  $T \leftarrow \text{InitializeTree}(x_{start}, y_{start}, \theta_{start}, x_{goal}, y_{goal})$ 
3 for  $i = 0, \dots, N$  do
4    $e \leftarrow \text{random}()$ 
5   if  $e < E$  or not  $T.\text{pathToGoalFound}()$  then
6     if  $e < E$  then // Sample free space
7        $x_{new}, y_{new} \leftarrow \text{sampleNavigableXY}()$ 
8     else // Sample near shortest path
9        $x_{new}, y_{new} \leftarrow \text{sampleNearShortestPath}()$ 
10       $x_{new}, y_{new} \leftarrow T.\text{project}(x_{new}, y_{new})$ 
11    else // Sample near current fastest path
12       $x_{new}, y_{new} \leftarrow T.\text{sampleNearFastestPath}()$ 
13       $p_{near}^{1 \dots N} \leftarrow T.\text{nearbyNodes}(x_{new}, y_{new})$ 
14       $p_{near}, \theta_{new} \leftarrow \text{pickParent}(x_{new}, y_{new}, p_{near}^{1 \dots N})$ 
15       $p_{new} \leftarrow T.\text{initializeNode}(x_{new}, y_{new}, \theta_{new})$ 
16       $T.\text{addEdge}(p_{near}, p_{new})$ 
17       $T.\text{rewire}(p_{new})$ 
18 return  $T.\text{fastestPath}()$ 

```

contains 994 episodes across 14 scenes. We divide the validation split into a val_1 and val_2 set by randomly selecting 7 out of the 14 scenes for each set (497 episodes in each set). After training each agent on the train split, we choose the policy weights for each agent based on their best SCT performance on val_1 . Each agent then uses these weights to perform the episodes of val_2 , and performance on val_2 is used to report evaluation metrics.

B. Action space

We use discrete action spaces for training both point-turn and unicycle agents. At each step, the policy is polled and outputs an action representing a pair of linear and angular velocities. The simulator has the agent assume its output velocities (unless it collides) until it is polled again. The policy is polled at 1 Hz. The agent terminates the episode by outputting zero linear and angular velocity. We constrain



Fig. 5: Discrete action space used by each agent. Actions are scaled by the agent’s maximum linear and angular velocities.

the linear velocity to be non-negative, as we have found that allowing the agent to move backwards hurts performance; the resulting behavior often moves backwards excessively and does not use information from the agent’s depth camera effectively, as the agent’s camera is forward-facing.

As shown in Figure 5, we use discrete action spaces with different amounts of actions to model different dynamics:

- Point-turn (2 lin. vel. + 3 ang. vel. - 1 overlap = 4 actions)
- Unicycle-6 (2 lin. vel. \times 3 ang. vel. = 6 actions)
- Unicycle-15 (3 lin. vel. \times 5 ang. vel. = 15 actions)

We set the maximum linear velocity to 0.25m/s and the maximum angular velocity to 10°/s. For the actions of the point-turn agent, this translates to 0.25m or 10° of displacement per step, which is used by other works such as [12], [16], [18], [20].

C. Observation space

The robot is equipped with an egocentric depth camera and an egomotion sensor. The robot’s depth camera provides single channel 256 \times 256 images where each pixel represents a depth of up to 10m. The egomotion sensor provides the agent with its current displacement and relative heading from its initial start position. This is used to calculate the relative heading and distance to the goal from the goal coordinates, which is then given to the robot’s policy.

D. Policy model

We use the neural net architecture described in [20] to model the policy of each of our agents. It has two main components; a visual encoder and a policy. The visual encoder is a convolutional neural network based on ResNet-50 [8], and takes the depth image as input. The policy is comprised of a 2-layer LSTM recurrent neural network with a 512-dimensional hidden state, taking the visual encoder’s features, the relative distance and heading to the goal, the previous action, and the previous hidden state as input (Figure 6). The goal distance and heading are represented as $[r, \cos(\theta), \sin(\theta)]$ to avoid the discontinuity at 180°, and the previous action is represented as a one-hot encoded vector, whose length depends on the agent’s action space. The final layer of the policy has either 4, 6, or 15 outputs (depending on the agent’s action space), which parameterize a categorical action distribution from which to sample from. The policy

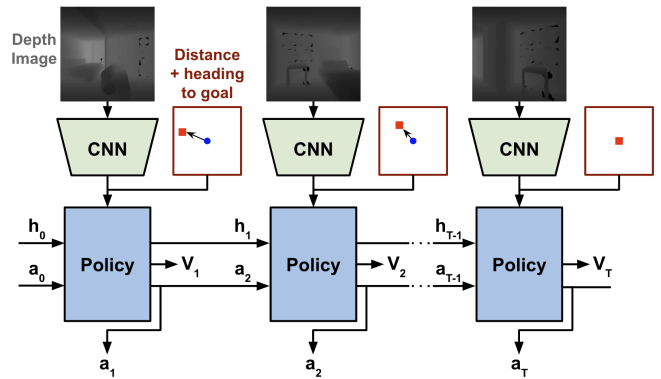


Fig. 6: Network architecture for our agents, comprised of a convolutional encoder and a recurrent policy. The agent receives an egocentric depth image, its current distance and heading relative to the goal point, and its previous action at each time step. It outputs an action and an estimate of the value function.

also has a critic head that outputs an estimate of the state’s value, which is used for reinforcement learning as described in [20].

E. Training and reward

We use Decentralized Distributed Proximal Policy Optimization by Wijnmans et al. [20] to train our agents using deep reinforcement learning, and use the same training hyperparameters. We train 8 copies of the agent running in parallel per GPU, using 8 GPUs for a total of 64 parallel workers, for 150M cumulative steps (\sim 240 GPU hours, or \sim 30 hours wall-clock).

We consider two different reward schemes. The first reward scheme is based on the one used by Wijnmans et al., which has shown to yield very good performance ($>96\%$ success, >0.91 SPL) using a point-turn agent. However, to make learning easier, it uses a shaped step reward that explicitly encourages the agent to follow the shortest path. Under such a reward scheme, the agent may not learn to take a path that deviates from the shortest path, even if it would allow the agent to reach the goal faster. Because we aim to evaluate the agent based on how well it approximates the fastest completion behavior with SCT, we thus consider a second reward scheme that “un-shapes” this step reward, by gradually reducing its shaped portion to 0 as training progresses.

Shaped reward scheme. In the reward scheme used by Wijnmans et al., the agent receives a shaped reward $r(a_t, s_t)$ at each time step based on geodesic distance, which is defined as the length of the shortest obstacle-free path between two points. We compute geodesic distance using A* on a navigation mesh provided by the simulator. When the agent reduces the geodesic distance from its position to the goal point, the change in geodesic is negative; thus, we multiply by -1 to produce a positive shaped reward if the agent reduces its shortest route to the goal. A slack penalty (-0.01) is added to discourage the robot from taking unnecessary actions. Lastly, a terminal reward r_T is given at the end of the episode if the goal has been reached. Though Wijnmans et al. scale the terminal reward by the agent’s SPL for the episode, we avoid

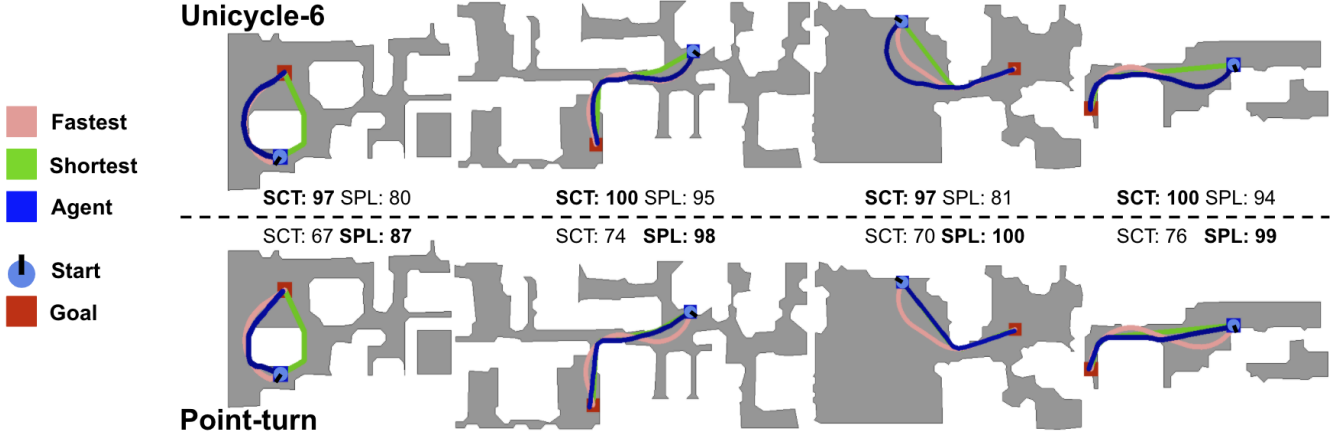


Fig. 7: The unicycle agents learn to leverage their dynamics to complete episodes faster than point-turn agents (higher SCT) by using smoother turns, which leads to longer paths and consequently lower SPL. Visualization is of performance on example episodes from val_2 , trained with the decaying reward.

this to refrain from explicitly penalizing agents for deviating from the shortest path. The full reward scheme is defined as:

$$r(a_t, s_t) = -\Delta_{geo_dist} - 0.01, \quad r_T = 2.5S \quad (3)$$

where S is 0 or 1 depending on whether the agent successfully completed the episode.

Decaying reward scheme. As previously mentioned, our aim with the second reward scheme is to prioritize successful completion speed. We keep the terminal reward the same, but slowly phase the shaped portion of the step reward to 0 as training progresses:

$$r_{decay}(a_t, s_t) = -\beta\Delta_{geo_dist} - 0.01, \quad r_T = 2.5S \quad (4)$$

where β starts with a value of 1, and linearly decays to 0 as the amount of elapsed training steps reaches a certain amount (80 million steps was chosen for our experiments). The goal of this reward scheme is to eventually provide the agent with only a negative constant step reward and a large positive constant terminal reward; the latter encourages the agent to complete the episode successfully, while the former encourages it to do it as quickly as possible. We do not immediately remove the geodesic portion of the reward at the start of training because the reward may be too sparse to learn optimal behaviors. Specifically, an agent early in training that has likely not learned to consistently reach the goal yet would be encouraged to simply terminate the episode immediately, such that it can avoid accruing the negative step reward. The geodesic term provides the positive reward signal necessary, even during the early stages of training, to lead the agent to the goal and teach it about the existence of a large positive terminal reward. Once the agent has learned that continuing the episode can lead to a large reward, we find that it continues to do so even as the geodesic term is completely eliminated.

V. RESULTS

We aim to address the following questions:

- 1) How effective are SCT and SPL for capturing the navigation efficiency of agents with a dynamics model more complex than a point-turn model?

- 2) How does un-shaping the reward function affect SCT and other metrics of performance across different dynamics?
- 3) Do our methods and results generalize to reality?

A. Evaluation details

After training, we set the agent’s behavior to be deterministic for both the val_1 and val_2 sets, meaning that the action with the highest likelihood in the output action distribution is always chosen. The agent does not receive any rewards, relying entirely on its depth camera and egomotion sensor as input to reach the goal. As detailed in Section IV-A, all reported evaluation metrics are based on performance on val_2 . All reported SCT, SPL, and success rate values are multiplied by 100 for readability. To calculate SCT, we use RRT*-Unicycle to approximate the fastest possible path time for a unicycle-cart model of dynamics constrained to a maximum linear velocity of 0.25m/s and an angular velocity of 10°/s. Although the fastest completion times of SCT are computed based on unicycle dynamics, calculating SCT for the point-turn agent in the same way as the unicycle agents allows us to properly evaluate how much faster or slower it is successfully completing the test episodes in comparison.

B. Path Length vs. Completion Time

The results of our experiments within simulation are shown in Table I. Average SCT and SPL values are highly correlated with success rate, as unsuccessful episodes will yield 0 SCT and 0 SPL. This complicates the comparison of agents with *varying success rates* in terms of average path length and completion time. We thus remove success rate as a variable by re-calculating average SCT and SPL using only the episodes that *all 6* agent variants (3 agents \times 2 reward schemes) succeeded in (totaling to 87% of val_2), dubbing them SCT^\cap and SPL^\cap . These values show that for successful episodes, the unicycle agents completed episodes much faster than the point-turn agent ($>88 SCT^\cap$ vs. $\sim 72 SCT^\cap$), but took paths that were not as short ($\sim 90 SPL^\cap$ vs. $\sim 95 SPL^\cap$).

TABLE I: Performance comparison between Point-turn, Unicycle-6, and Unicycle-15 agents using the shaped and decaying reward schemes. Mean and 95% confidence interval are reported.

Agent	Decay	SCT \uparrow	SPL \uparrow	Succ. \uparrow	SCT $^{\square}$ \uparrow	SPL $^{\square}$ \uparrow
PT	-	65.82 \pm 1.95	88.41 \pm 2.18	94.16 \pm 2.06	71.26 \pm 1.33	95.11 \pm 0.92
	\checkmark	65.99 \pm 2.07	87.91 \pm 2.28	93.96 \pm 2.09	71.75 \pm 1.45	94.90 \pm 1.08
U-6	-	82.15 \pm 2.56	82.83 \pm 2.42	93.16 \pm 2.22	90.41 \pm 1.59	90.70 \pm 1.35
	\checkmark	88.92 \pm 1.84	88.07 \pm 1.61	98.79 \pm 0.96	93.66 \pm 1.19	91.64 \pm 1.02
U-15	-	81.43 \pm 2.53	83.31 \pm 2.26	94.37 \pm 2.03	88.37 \pm 1.82	89.69 \pm 1.41
	\checkmark	86.01 \pm 2.10	85.37 \pm 1.87	95.77 \pm 1.77	91.09 \pm 1.56	89.66 \pm 1.34

This finding supports our argument that SPL may not be a suitable metric for robots with more realistic dynamics. For a model of dynamics such as the unicycle-cart, it is unclear what value of SPL is ‘optimal’. On one hand, the maximum value of 100 SPL would restrict a unicycle agent’s behavior to only move with its point-turn actions. On the other, targeting a lower SPL value that better corresponds to a learned behavior that more efficiently reaches the goal by leveraging non-point-turn dynamics is difficult, as it would be hard to determine an appropriate target SPL. Since SCT uses fastest completion times conditioned explicitly on the agent’s dynamics as a baseline, it is a better metric for directly measuring how well the agent’s learned behavior leverages the extents of its dynamics for navigation.

C. Reward scheme comparison

Table I shows that relative to the shaped reward scheme, our decaying reward scheme significantly increased the average SCT of the unicycle agents (82.2 \rightarrow 88.9 SCT for Unicycle-6, 81.4 \rightarrow 86.0 SCT for Unicycle-15), even with increases in success rate (93.2% \rightarrow 98.8% for Unicycle-6, 94.4% \rightarrow 95.8% for Unicycle-15) removed as a variable (90.4 \rightarrow 93.7 SCT $^{\square}$ for Unicycle-6, 88.4 \rightarrow 91.1 SCT $^{\square}$ for Unicycle-15). On the other hand, the performance of the point-turn agent did not change much between reward schemes (<0.5 change for SCT/SPL/SCT $^{\square}$ /SPL $^{\square}$, 0.2% change in success rate). We believe that this is because the original reward scheme encourages the agent to make shorter paths to the goal, while the decaying reward scheme encourages faster ones; these paths are not much different for a point-turn agent, so the learned behavior would not change much. Our results show that by gradually removing explicit encouragement of following the shortest path, the unicycle agents learn to better leverage the capabilities of their dynamics to complete the episodes faster and succeed more frequently.

D. Real-world experiments

We conducted real-world experiments in an indoor apartment to test whether point-turn and unicycle policies trained in simulation using our decaying reward scheme can generalize to reality. We used the LoCoBot mobile robot platform [7], equipped with a Hokuyo UTM-30LX LiDAR. The raw LiDAR data is *not* processed by the robot’s policy, and is used only to provide local egomotion estimates (relative heading and location from start position). The robot is also equipped with an Intel D435 depth camera to provide the policy with depth imagery. To calculate the completion time

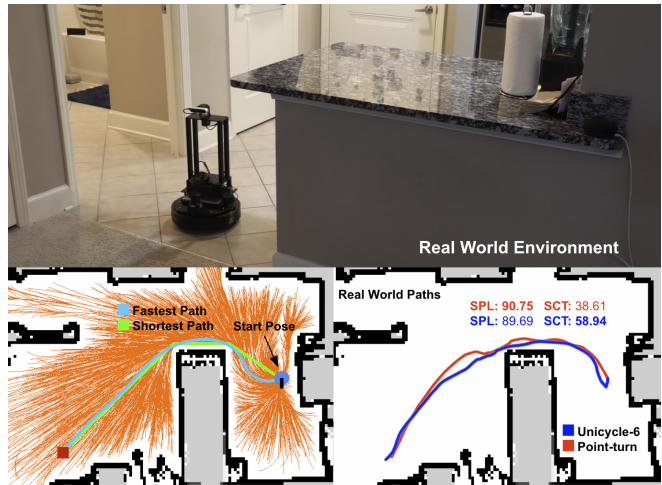


Fig. 8: *Top*: The LoCoBot mobile robot and apartment used in real-world experiments. *Bottom left*: RRT*-Unicycle and RRT*-Vanilla used on a 2D LiDAR scan to find the fastest completion time and shortest path length, respectively. *Bottom right*: Example runs from the two agents on Ep. ID #1 shown on the right.

TABLE II: Results from reality experiments. Each episode is the average of three independent trials. The Unicycle-6 agent achieves a higher SCT and faster completion time, while the Point-turn agent achieves a higher SPL and shorter distance travelled.

Ep. ID	Agent	SCT \uparrow	SPL \uparrow	Time (sec.) \downarrow	Agent Dist. (m) \downarrow
1	PT	36.80 \pm 2.13	88.37 \pm 2.82	61.00 \pm 3.61	5.59 \pm 0.18
	U-6	54.01 \pm 4.62	85.46 \pm 3.99	41.67 \pm 3.51	5.78 \pm 0.27
2	PT	37.99 \pm 5.60	93.79 \pm 0.81	86.00 \pm 13.86	8.03 \pm 0.07
	U-6	59.96 \pm 6.05	92.87 \pm 3.49	54.00 \pm 5.57	8.11 \pm 0.31
Avg.	PT	37.39 \pm 3.84	91.08 \pm 3.50	73.50 \pm 16.42	6.81 \pm 1.34
	U-6	56.98 \pm 5.81	89.16 \pm 5.27	47.83 \pm 7.94	6.95 \pm 1.30

of the fastest path, we use RRT*-Unicycle on the floorplan generated by the Hector SLAM algorithm [11] using the LiDAR. Neither the calculated path nor the floorplan is provided to the agent, and is used for performance evaluation purposes only. Because a navigation mesh of the apartment is not available, we simply use a vanilla version of RRT* to find the shortest path and its length for SPL calculation. These paths are visualized in Figure 8.

We tested two navigation agents, Point-turn and Unicycle-6. We increase their maximum angular speed during training and testing to 30 $^{\circ}$ /s, as we have found that commanding 10 $^{\circ}$ /s is too slow, often resulting in little to no movement on the real LoCoBot. This translates to 30 $^{\circ}$ per step (as the policy is polled at 1 Hz), which has shown to work for visual navigation using the LoCoBot in [9]. We do not test with the Unicycle-15 agent, as its action space would contain angular speeds of 15 $^{\circ}$ /s, which we also found to be too slow. Following the recommendations from Kadian et al. [9], we train these agents with Habitat’s sliding setting set to ‘off’ in the simulator settings to help minimize collisions with obstacles. Sliding is a setting which allows the agents to ‘slide’ against walls and obstacles upon collisions, a behavior often used in video games and other physics-based simulators. Kadian et al. showed that agents trained with

this setting turned on often abuse these collision dynamics in simulation to reach the goal, and collide much more often when deployed in the real world. When sliding is off, the agent does not move upon collision. Otherwise, training and checkpoint selection was done in the same manner as our 10⁰/s agent experiments reported in the previous subsections.

Each agent was tested on two different episodes, each with 3 independent trials, for a total of 12 real-world runs. A top-down view of one episode is shown in Figure 8. Note that our agents were trained solely in simulation, and were not fine-tuned to the evaluation environment.

Our results in Table II show that the Unicycle-6 agent reached the goal significantly faster than the Point-turn agent on average (48 sec vs. 74 sec), while taking a slightly longer path (7.0 m vs 6.8 m). Additionally, we show that SCT accurately highlights the efficiency of the faster agent, as the Unicycle-6 agent achieves a higher average SCT than the Point-turn agent (57 SCT vs. 37 SCT), while considering SPL alone may make the Unicycle-6 agent seem less efficient than the Point-turn agent (89 SPL vs 91 SPL).

VI. CONCLUSION

In this work, we propose a new metric, Success by Completion Time. In contrast to other navigation performance metrics like SPL, SCT explicitly considers the dynamics model of the agent’s embodiment. This allows SCT to capture how well an agent’s learned navigation behavior approximates the ideal behavior afforded by its dynamics, where for SCT, the ideal behavior is defined as one that results in the fastest successful completion times. We also present our RRT*-Unicycle algorithm for approximating the fastest path in an environment with obstacles for an agent with unicycle dynamics, and show that it can be used to evaluate the navigation performance of agents both in simulation and in reality. Using deep reinforcement learning, we train unicycle agents that complete episodes much faster than point-turn agents not by having increased speed, but by learning to leverage their dynamics. We also propose a decaying reward scheme that encourages this behavior and significantly improves SCT performance. Lastly, we show that agents trained solely in simulation with our decaying reward scheme can generalize to reality, with the unicycle agent completing episodes faster than the point-turn agent.

There are multiple potential future directions of research available to extend our work. Because discrete action spaces are only able to represent a subset of the full space of actions of a dynamics model, it limits the agent’s ability to reach the fastest behavior of its dynamics. Thus, research in training unicycle agents with a continuous action space may lead to agents that surpass our agents and achieve near-optimal SCT performance. Another promising direction would be to investigate other models of dynamics, such as car-like robots (that cannot pivot) or legged robots, and use SCT as a better metric for navigation performance than SPL.

REFERENCES

- [1] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir Roshan Zamir. On evaluation of embodied navigation agents. *CoRR*, abs/1807.06757, 2018. 2
- [2] Dhruv Batra, Aaron Gokaslan, Aniruddha Kembhavi, Oleksandr Maksymets, Roozbeh Mottaghi, Manolis Savva, Alexander Toshev, and Erik Wijmans. Objectnav revisited: On evaluation of embodied agents navigating to objects, 2020. 1, 2
- [3] Devendra Singh Chaplot, Dhiraj Gandhi, Abhinav Gupta, and Ruslan Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, page 33, 2020. 1, 2
- [4] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *International Conference on Learning Representations (ICLR)*, 2020. 1
- [5] Changan Chen, Ziad Al-Halah, and Kristen Grauman. Semantic audio-visual navigation, 2020. 2
- [6] Saurabh Satish Desai and Stefan Lee. Auxiliary tasks for efficient learning of point-goal navigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 717–725, January 2021. 2
- [7] Abhinav Gupta, Adithyavairavan Murali, Dhiraj Gandhi, and Lerrel Pinto. Robot learning in homes: Improving generalization and reducing dataset bias. *CoRR*, abs/1807.07049, 2018. 2, 7
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 5
- [9] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Are we making real progress in simulated environments? measuring the sim2real gap in embodied visual navigation. *CoRR*, abs/1912.06321, 2019. 1, 7
- [10] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011. 3
- [11] S. Kohlbrecher, J. Meyer, O. von Stryk, and U. Klingauf. A flexible and scalable slam system with full 3d motion estimation. In *SSRR*. IEEE, November 2011. 7
- [12] Shangda Li, Devendra Singh Chaplot, Yao-Hung Hubert Tsai, Yue Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. Unsupervised domain adaptation for visual navigation, 2020. 1, 2, 5
- [13] Chang-bae Moon, Woojin Chung, and Nakju Lett Doh. Observation likelihood model design and failure recovery scheme toward reliable localization of mobile robots. *International Journal of Advanced Robotic Systems*, 7(4):24, 2010. 1
- [14] Jauwairia Nasir, Fahad Islam, Usman Malik, Yasar Ayaz, Osman Hasan, Mushtaq Khan, and Mannan Saeed Muhammad. Rrt*-smart: A rapid convergence implementation of rrt*. *International Journal of Advanced Robotic Systems*, 10(7):299, 2013. 3
- [15] Iram Noreen, Amna Khan, Zulfiqar Habib, et al. Optimal path planning using rrt* based approaches: a survey and future directions. *Int. J. Adv. Comput. Sci. Appl*, 7(11):97–107, 2016. 3
- [16] Marco Rosano, Antonino Furnari, Luigi Gulino, and Giovanni Maria Farinella. On embodied visual navigation in real environments through habitat, 2020. 2, 5
- [17] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A platform for embodied AI research. *CoRR*, abs/1904.01201, 2019. 1, 4
- [18] Tianqi Tang, Xin Yu, Xuanyi Dong, and Yi Yang. Auto-navigator: Decoupled neural architecture search for visual navigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3743–3752, January 2021. 2, 5
- [19] Dustin J Webb and Jur Van Den Berg. Kinodynamic rrt*: Asymptotically optimal motion planning for robots with linear dynamics. In *2013 IEEE International Conference on Robotics and Automation*, pages 5054–5061. IEEE, 2013. 3
- [20] Erik Wijmans, Abhishek Kadian, Ari S. Morcos, Stefan Lee, Irfan Essa, D. Parikh, M. Savva, and Dhruv Batra. Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames. In *ICLR*, 2020. 1, 2, 5
- [21] Fei Xia, Amir R. Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 4

[1] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra