Reinforcement Learning-based Visual Navigation with Information-Theoretic Regularization

Qiaoyun Wu¹, Kai Xu², Jun Wang¹, Mingliang Xu³, Xiaoxi Gong¹, and Dinesh Manocha⁴

Abstract-To enhance the cross-target and cross-scene generalization of target-driven visual navigation based on deep reinforcement learning (RL), we introduce an information-theoretic regularization term into the RL objective. The regularization maximizes the mutual information between navigation actions and visual observation transforms of an agent, thus promoting more informed navigation decisions. This way, the agent models the action-observation dynamics by learning a variational generative model. Based on the model, the agent generates (imagines) the next observation from its current observation and navigation target. This way, the agent learns to understand the causality between navigation actions and the changes in its observations, which allows the agent to predict the next action for navigation by comparing the current and the imagined next observations. Cross-target and cross-scene evaluations on the AI2-THOR framework show that our method attains at least 10%improvement of average success rate over some state-of-the-art models. We further evaluate our model in two real-world settings: navigation in unseen indoor scenes from a discrete Active Vision Dataset (AVD) and continuous real-world environments with a TurtleBot. We demonstrate that our navigation model is able to successfully achieve navigation tasks in these scenarios¹.

Index Terms—Visual-Based Navigation, Reinforcement Learning, Model Learning for Robot Control.

I. INTRODUCTION

VISUAL navigation is one of the basic components necessary for an autonomous agent to perform a variety of tasks in complex environments. This component can be characterized as the ability of an agent to understand its surrounding environments and navigate efficiently and safely to a designated target solely based on input from on-board

Manuscript received: August, 21, 2020; Revised October, 30, 2020; Accepted December, 15, 2020

This paper was recommended for publication by Editor in Chief Allison Okamura and Editor Eric Marchand upon evaluation of the Associate Editor and Reviewers' comments. This work was supported in part by the grants from National Key Research and Development Program of China (2020YFB2010702, 2019YFB1707504, 2018AAA0102200), National Natural Science Foundation of China (61902419, 61772267, 61622212, 61532003), Aeronautical Science Foundation of China (2019ZE052008), and the National Science Foundation of Jiangsu Province (BK20190016). Corresponding author: Kai Xu (kevin.kai.xu@gmail.com). Corresponding author: Jun Wang (wjun@nuaa.edu.cn).

¹Q. Wu, J. Wang, and X. Gong are with the College of Mechanical and Electrical Engineering, Nanjing University of Aeronautics and Astronautics, 210016 China.

²K. Xu is with the School of Computer Science, National University of Defense Technology, 410003 China.

³M. Xu is with the Zhengzhou University, 450052 China.

⁴D. Manocha is with the Department of Computer Science, the University of Maryland, College Park, MD 20742 USA.

This article has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the authors.

Digital Object Identifier (DOI): see top of this page.

¹https://github.com/wqynew/RL-based-navigation

visual sensors [1], [2]. This encompasses two key points. First, the agent should be able to analyze and infer the aspects most relevant to the target from the current observation to guide the decision. Second, the agent should understand the correlation and causality between navigation actions and the changes in its observation of the surroundings.

Recently, there has been an increased interest in mapless visual navigation approaches where the agent neither relies on the prior knowledge of the environment, nor performs online mapping. Instead, it predicts navigation actions directly from observational pixels thanks to end-to-end deep learning, e.g., Imitation Learning (IL) [3] and Deep Reinforcement Learning [4], [5]. Despite significant progress in visual navigation, the generalization to novel targets and unseen scenes is still a fundamental challenge. In this work, we focus on visual navigation, driven by targets represented by an image, with both cross-scene and cross-target generalization.

To achieve visual navigation, we propose enhancing a Deep Reinforcement Learning approach (e.g. A3C [1]) with an information-theoretic regularization. We introduce the regularization into the RL objective to guide the agent in a more informative search for its navigation actions. In particular, the regularization maximizes the mutual information between the action and the next visual observation given the current visual observation of the agent. This way, the agent models the action-observation dynamics and learns to understand the causality between navigation actions and the changes in its observations, thus making more informed decisions.

The maximization is, however, intractable due to the unknown next visual observation at each time step. Inspired by our previous work [6], which presents a variational Bayesian model (NeoNav) for supervised navigation learning², we introduce a variational auto-encoder (VAE) model, which generates (imagines) the next observation based on the current observation and the target view. We regularize the latent space of the VAE through the action-observation dynamics. The agent then learns to predict the next action based on the current and the *imagined* observations. Consequently, the agent essentially builds a connection between the current observation and the target to infer the most relevant part for navigation and makes decisions based on the causality between navigation actions and observational changes.

There are several works on introducing informationtheoretic regularization to RL [7], [8]. Most of them strive to maximize the entropy of the policy to encourage exploration or to make the policy more stochastic for better robustness. A

²https://github.com/wqynew/NeoNav

specifically related work is [9], which devises a similar mutual information maximization as an internal reward for learning an intrinsically motivated agent. In contrast to their work, we use mutual information maximization as a regularization of the objective and learn a generative model of the actionobservation dynamics. To our knowledge, our method is the first to use information-theoretic regularization to guide the learning of generalizable visual navigation.

In summary, our contributions are as follows: (1) We present a novel method of incorporating supervision into an RL framework (A3C) by introducing an information-theoretic regularization for improving the sample efficiency in targetdriven visual navigation policy learning. (2) We propose a visual navigation model that builds the causality between robot actions and visual transformations, and the connection between visual observations and navigation targets. This significantly improves robot navigation performance in unseen environments with novel targets. We conduct extensive evaluations on datasets from both synthetic and real-world scenes, including AI2-THOR [1] and AVD [10]. Our model outperforms some state-of-the-art methods significantly (e.g., at least 10% higher success rate for both cross-target and cross-scene evaluation on AI2-THOR). Furthermore, we show that our model, trained on the discrete household dataset (e.g., AVD) and deployed on a Turtlebot, can transfer to realistic public scenes and exhibit robustness towards the target type and the scene layout.

II. RELATED WORKS

Autonomous navigation in an unknown environment is one of the core problems in mobile robotics and it has been extensively studied. In this section, we provide a brief overview of some relevant works.

Reinforcement Learning. Recently, a growing number of methods have been reported for RL-based navigation [11], [12], [13], [14]. Gupta et al. [15] present an end-to-end architecture to jointly train mapping and planning for navigation in novel scenes with the perfect odometry available assumption. Savinov et al. [16] propose the use of topological graphs for navigation tasks and require several minutes of footage before navigating in an unseen scene. Wei et al. [17] integrate semantic and functional priors to improve navigation performance and can generalize to unseen scenes and objects. Xie et al. [18] propose using a few snapshots of the environment combined with directional guidance to help execute navigation tasks. Hirose et al. [19] introduce a learning agent that can follow a demonstrated path. The path consists of raw image sequences when navigating in a new environment, which largely discounts the practicality.

Combined Learning Methods. Methods combining the advantages of IL and RL have become popular [20]. These works provide suitable expert demonstrations to mitigate the low RL sample efficiency problem. Ho et al. [21] exploit a generative adversarial model to fit distributions of states and actions defining expert behavior. They learn a policy from supplied data and hence avoid the costly expense of RL. [22], [23] share the same idea of learning from multiple teachers. Li et al. [22] discard bad maneuvers by using a

reward-based online evaluation of the teachers during training. Muller et al. [23] use a DNN to fuse multiple controllers and learn an optimized controller. Target-driven navigation in static environments is different from the problems above due to the easy acquisition of the optimal expert (the shortest path). Hence, there is no need to consider the bad demonstrations. We learn to maximize an expected long-term return provided by environments. On the other hand, we add an intermediate process to the navigation policy (the generation of the future observation) and predict an action based on the difference between the current and the future observations. This makes a more effective and generalizable navigation model.

Information Gain-based Approaches. Information gainbased strategies have been applied to a variety of robotics problems involving planning and control. They have been used to study optimal sensor placement and motion coordination for a target-tracking task [24], derive an information-theoretic metric as a new visual feature for visual servoing [25], optimize an information-theoretic objective to improve the informativeness of both local motion primitives and global plans for mapping [26], facilitate RL to compute good trajectories for scene exploration [7], and generate intrinsic reward to learn an exploration policy [27]. There are differences in the way mutual information is used in these applications. However, information gain-based strategies have not been applied to target-driven visual navigation, the goal of which is to navigate from a random location in a scene to a specified target.

Dynamics Model Learning. There is extensive literature on learning a dynamics model and using this model to train a policy. Most notable among these is the work from [28], which proposes that the Imagination-Augmented Agent learns approximate environment models before outputting the action policy. Ha et al. [29] incorporate a generative recurrent model into reinforcement learning to predict the future given the past in an unsupervised manner. However, the goal information is hard-coded in these neural networks and the experimental environments are generally simple and fully observed, leading to poor generalization to complex, high-dimensional tasks with unseen targets in partially observed scenes. Pathak et al. [30] learn an inverse dynamics model based on the demonstrated trajectory way-points from an expert and demonstrate navigation in previously unseen office environments with a TurtleBot. We learn a dynamics model by maximizing the mutual information between the action and the adjacent observation pair, which can achieve better navigation generalization.

III. TARGET-DRIVEN VISUAL NAVIGATION

In this section, we begin by outlining the target-driven visual navigation task. We then present our network, which combines an information-theoretic regularization with deep reinforcement learning for this task.

A. Navigation Task Setup

We focus on learning a policy for navigating an agent from its current location to a target in an unknown scene using only visual observations. Our problem is: given a target image g, at each time step t, the agent receives as input an observation x_t of an environment to predict an action a_t that will navigate the robot to the viewpoint where g is taken.

Datasets. We conduct our experiments on the AI2-THOR, AVD, and some real-world scenarios. AI2-THOR consists of 120 synthetic scenes in four categories: kitchen, living room, bedroom, and bathroom. Each category includes 30 scenes, 20 of which are used for training, 5 for validating, and 5 for testing, in accordance with [2]. AVD contains 14 different households, 8 of which are used for training, 3 for validating, and 3 for testing, as in [10]. We further transfer the learned policies from AVD to some real-world public scenes based on a robotic platform (e.g., TurtleBot); these scenes have never been encountered before.

Observations. In contrast to [1], which stacks four history frames as current inputs at each time step, we utilize four views (RGB images by default) with evenly distributed azimuth angles at each location for current observation x_t . The resolution of each view is 300 * 300.

Targets. The navigation target is specified by an RGB image, which contains a goal object. Please refer to [31] for the training and testing goal objects. Our model learns to analyze the relationship between the current observation and the target image, and hence we can show generalization to novel targets and scenes that the agent has not previously encountered.

Actions. Each scene in our dataset is discretized into a grid-world navigation graph. The agent acts on these graphs and its action space is determined by the connectivity structures of these graphs as a discrete set: $\mathcal{A} =$ {move forward; move back; move left; move right; rotate ccw; rotate cw; stop}, as defined in [10]. The discrete graph makes it easy to acquire a shortest action path for a target-driven navigation task (e.g., using A-star algorithm). In this work, we will show how to incorporate the shortest paths during training to learn a navigation controller.

Rewards. Our purpose during policy training is to minimize the length of the trajectory to the navigation targets. Therefore, reaching the target is assigned a high reward value 10.0 and we penalize each step with a small negative reward -0.01. To avoid collision, we design a penalty of -0.2 when obstacles are hit during run-time. In addition, we consider the geodesic distance to the goal at each time step, $Geo(x_t; g)$, as in [5], and reformulate the reward as:

$$r_{t} = \begin{cases} -0.01 & \text{if } t = 0 \\ +10.0 & \text{elif } succeed \\ -0.2 & \text{elif } collide \\ Geo(x_{t-1}, g) - Geo(x_{t}, g) - 0.01 & \text{otherwise.} \end{cases}$$
(1)

Success measure. In our setting, the agent runs up to 100 steps, unless a stop action is issued or a task is successful. The task is considered successful if the agent predicts a stop action, the goal object is in the field of the current front-view, and the distance between the current location and the target view location is within a threshold (e.g., 0.5m for the AI2-THOR simulator and 1.0m for two real-world settings).

B. Information-Theoretic Regularization

We formulate the target-driven visual navigation using a deep reinforcement learning framework (TD-A3C). At each

time step t, the network takes a current observation x_t and a navigation target g as inputs and finally outputs an action distribution $\pi(x_t, g)$ and a scalar $v(x_t, g)$. We choose action a_t from the policy $\pi(x_t, g)$, and $v(x_t, g)$ is the value of the current policy. This network can be updated by minimizing a traditional RL navigation loss as in [1]. Figure 1(a) shows the interaction between the agent and an environment. However, achieving strong results with one single policy network for all training scenes is difficult, since the agent is very sensitive to the RL reward function and requires extensive training time. In addition, [1] does not consider generalization to previously unseen environments.



Fig. 1. Target-driven navigation flow diagram showing how agents interact with the environment. (a) Traditional RL agent (in the orange square). (b) Our agent (in the blue square). Layer parameters in a gray square are shared by a policy network and a value network. The blue represents the difference between the two flow diagrams. We propose using an information-theoretic regularization to facilitate the traditional RL learning, which requires more information from the environment, e.g., x_{t+1}^{gt} and a_t^{gt} .

In order to address the shortcoming above, we revisit Shannon's mutual information (MI) measure to further reduce the uncertainty in navigation action decisions when a visual observation is given. Let x_t denote the current observation, x_{t+1} denote the next observation, and a_t be the relative action between the two observations. We observe that an agent always abides by a task-independent information-theoretic regularization: there should be high mutual information between the action a_t and the next observation x_{t+1} given the current observation x_t . The mutual information $I(x_{t+1}, a_t | x_t)$ is defined as:

$$\begin{split} I(a_t, x_{t+1} | x_t) &= H(a_t | x_t) - H(a_t | x_{t+1}, x_t) \\ &= \iint p(a_t, x_{t+1} | x_t) \log p(a_t | x_t, x_{t+1}) da_t dx_{t+1} + H(a_t) \\ &\geq \iint p(x_{t+1} | x_t, a_t) p(a_t | x_t) \log p(a_t | x_t, x_{t+1}) da_t dx_{t+1} \\ &= \iiint p(x_{t+1} | z) p(z | x_t, a_t) p(a_t) \log p(a_t | x_t, x_{t+1}) da_t dx_{t+1} dz \\ &= E_{x_{t+1} \sim p(x_{t+1} | z)} [E_{z \sim p(z | x_t, a_t)} [E_{a_t \sim p(a_t)} [\log p(a_t | x_t, x_{t+1})]] \end{split}$$

In this setting, we suggest that the action a_t of an agent is unrelated to its current visual observation x_t , but in connection with x_t only if the next observation x_{t+1} or a navigation goal g is given. Thus, we have $p(a_t|x_t)$ identically equal to $p(a_t)$, i.e. $H(a_t|x_t)$ is identically equal to $H(a_t)$. This is different from traditional learning-based visual navigation methods, which tackle individual tasks in isolation, where the goal information is hard-coded in the neural networks and corresponding state descriptions [32]. Thus, $p(a_t|x_t) \neq p(a_t)$ and these present poor generalization to unexplored targets. In addition, our action space is the deterministic discrete set \mathcal{A} , i.e. $a_t \in \mathcal{A}$. Hence, we assume $a_t \sim Cat(1/C)$, where $Cat(\cdot)$ is a categorical distribution and C is the cardinality of \mathcal{A} . $H(a_t) \geq 0$ is a constant. This regularization provides a wellgrounded action-observation dynamic model $p(x_{t+1}|x_t, a_t)$ and describes the causality between navigation actions and observational changes $p(a_t|x_t, x_{t+1})$. An agent that seeks to maximize this value will gain a compelling understanding of the dynamics and the causality. This intuition leads us to incorporate the regularization into our navigation learning.

We propose adapting the task-independent regularization above by incorporating some supervision to help learn a strong target-driven visual navigation model, a special case of the lower bound in Equation 2. The supervision is from the shortest paths of target-driven navigation tasks. Specifically, at each time step, given the current observation and the target, the optimal next observation x_{t+1}^{gt} and relative action a_t^{gt} are provided as ground truth, see Figure 1(b).

To maximize the lower bound, we first assume the next observation x_{t+1} and the ground truth action a_t^{gt} are given and thus the lower bound is converted to a predictive control term as $E_1 = E_{a_t^{g_t} \sim p(a_t^{g_t})}[\log p(a_t|x_t, x_{t+1})]$. Subsequently, we generate the next observation x_{t+1} by a generative module $z \sim p(z|x_t, a_t), x_{t+1} \sim p(x_{t+1}|z)$, which is most related to the navigation task. Hence, we use the ground truth action a_t^{gt} to guide the generation: $z \sim p(z|x_t, a_t^{gt}), x_{t+1} \sim p(x_{t+1}|z)$. The ground truth x_{t+1}^{gt} is used to help update the generation module though a reconstruction term $E_2 = ||x_{t+1}^{gt} - x_{t+1}||_2$. In addition, considering that a_t^{gt} is unknown a priori during real navigation and is inherently determined by the navigation target g, we design the distribution $q(z|x_t, g)$ to approximate the distribution $p(z|x_t, a_t^{gt})$, which formulates a Kullback-Leibler (KL) divergence term as $E_3 = \mathcal{KL}[q(z|x_t, g)||p(z|a_t^{gt}, x_t)].$ $z \sim p(z|x_t, a_t^{gt}), x_{t+1} \sim p(x_{t+1}|z) \text{ and } z \sim q(z|x_t, g)$ constitute our variational auto-encoder module. Overall, we obtain a variational objective function as:

$$\mathcal{J}(x_t, g) = \alpha E_1 - \beta E_2 - \gamma E_3 \tag{3}$$

The hyper-parameter (α, β, γ) tunes the relative importance of the three terms: predictive control, reconstruction, and KL.

C. Regularized Navigation Model

The key idea in reinforcement learning for navigation is finding a policy $\pi(x_t, g)$ that can maximize expected future return. Within our regularized navigation framework, along with the environment reward, our agent puts a great deal of weight on the ability to understand the action-observation dynamics and the causality between actions and observational changes. This changes the RL problem to:

$$\pi^* = \arg \max_{\pi} E[\sum_{t=0}^{\infty} \tau^t r_t + \mathcal{J}(x_t, g)]$$
(4)

where r_t is a reactive reward provided by the environment at each time step and $\tau \in (0, 1]$ is a discount factor. Corresponding to the objective, the actor-critic structure in our regularized A3C framework is developed as Figure 2.

Policy Network. The inputs to the policy network are the multi-view images x_t and the target image g at each time step t. The network first learns to reason about some important information from the current observation based on the target,

which is then used to generate the next expected observation. This process is supervised by the action-observation dynamics $p(x_{t+1}|x_t, a_t^{gt})$ and the ground truth next observation x_{t+1}^{gt} . Information from the generated observation and the current observation is fused to form a joint representation, which is passed through the predictive control layer to predict the navigation action.



Fig. 2. Model overview. Our model integrates an information-theoretic regularization into an RL framework to constrain the intermediate process of the navigation policy. During training, our network is supervised by the environment reward r_t , the shortest path of the task in the form of the ground truth action a_t^{gt} , and the ground truth next observation x_{t+1}^{gt} . The parameters are updated by four loss terms: the reconstruction, the KL, the predictive control and the value. The first three terms in blue are introduced by the information-theoretic regularization. At test time, the parameters are fixed and our network takes the current observation and the target as inputs to generate the future state. Then it predicts the action based on the future and the current states. Layer parameters in the green squares are shared.

In addition, we investigate two techniques to improve the training performance. First, we find that when the previous action a_{t-1} is provided, the agent is less likely to move or rotate back and forth in a scene. This is reasonable since the ground truth action has no chance to contradict the previous action (e.g., move forward vs. move backward). Second, we apply a CNN module f to derive a state representation from an image and hence get the current state $f(x_t)$, the ground truth next state $f(x_{t+1}^{gt})$, and the goal state f(g). Please refer to [31] for the structure. We do not directly generate the next observation x_{t+1} . We generate the state representation, denoted as $s_{t+1} \sim p(s_{t+1}|z)$ and use this to compute the reconstruction loss and predict the navigation action. To avoid confusion, we will still use the description of generating the next observation given below. This simplification reduces the network parameters and the computational cost. As a result, our navigation policy is updated by:

$$\mathcal{L}_{p} = \alpha E_{a_{t}^{gt} \sim p(a_{t}^{gt})} [-\log p(a_{t}|f(x_{t}), s_{t+1}, a_{t-1})] + \beta ||s_{t+1} - f(x_{t+1}^{gt})||_{2}$$
(5)
+ $\gamma \mathcal{K} \mathcal{L}[q(z|f(x_{t}), f(g))||p(z|a_{t}^{gt}, f(x_{t}))]$

Value Network. We learn a value function from the penultimate connected layer of our policy $\pi(x_t, g)$, which represents the value of the current policy at the current navigation task, denoted as $v(x_t, g)$. This is associated with a value loss $\mathcal{L}_v = E_{x_t,r_t}[(R_t - v(x_t, g))^2]$, where R_t is the discounted accumulative reward defined by $R_t = \sum_{i=0}^{T-t} \tau^i r_{t+i} + v(x_{T+1}, g)$. Unlike previous work in [1] which directly uses the value $v(x_t, g)$ (embodied as a TD error) to help update the navigation policy, see Figure 1(a), our value term \mathcal{L}_v merely affects the shared layers of the policy in Figure 1(b). Hence, our value network functions as an auxiliary task, and we will show that this cooperation consistently outperforms the baseline without it in Section IV-B.

Therefore, the overall loss function is $\mathcal{L} = \mathcal{L}_p + \omega \mathcal{L}_v$, where the hyper-parameter is empirically set as ($\alpha = 1.0$, $\beta = 0.01, \gamma = 0.0001, \omega = 0.5$) throughout our experiments. *Please refer to [31] for more details.*

At test time, three modules $z \sim q(z|x_t, g), x_{t+1} \sim p(x_{t+1}|z)$, and $a_t \sim p(a_t|x_t, x_{t+1}, a_{t-1})$ constitute our controller for the agent to predict the next action given the current observation, the target view, and the previous action. The controller can navigate robots in unseen scenes, the environment maps (graphs) of which are not known.

IV. IMPLEMENTATION AND PERFORMANCE

Our objective is to improve the cross-target and cross-scene generalization of target-driven navigation. In this section, we evaluate our model compared to baselines based on standard deep RL models and/or traditional imitation learning. We also provide ablation results to gain insight into how performance is affected by changing the structures.

A. Baselines and Ablations

We compare our method with the following visual navigation models: (1) TD-A3C is the target-driven visual navigation model from [1] and is trained using standard reinforcement learning. This was originally designed for scene-specific navigation and it is difficult to achieve strong results with one single policy network for all training scenes. We assist the policy learning by using the reward function and previous action as ours. (2) TD-A3C(BC) is a variation of the TD-A3C. It is trained using behavioral cloning (BC). Both the CNN module and the input are the same as ours. The main difference from our method is how the supervision is exploited. (3) Gated-LSTM-A3C(BC) is an LSTM-based variant of A3C model adapted from [4], which is trained with BC and provided with the previous action. The goal is specified as an image, and the model is also provided with the same multi-view images as in ours. (4) **GSP** is a goal-conditioned skill policy in [30], which generates the next observation as an auxiliary task rather than using the generation for navigation control. We reimplement the work³ and train it on our datasets. (5) SAVN proposes a self-adaptive visual navigation model [2], which shows strong results on novel scene adaption on AI2-THOR. It does not, however, support adaptation to novel targets. (6) **TD-Semantic** is a navigation model from [10]. The method predicts the cost of an action, which is supervised by shortest paths of navigation tasks. (7) NeoNav is our previous work [6]. (8) Ours-NoVal is a variant of our method, which does not consider the value of our navigation policy. (9) Ours-FroView is a variant of our method and takes the current front-view to generate the future observation. (10) Ours-NoGen is a variation of our model that predicts x_{t+1} directly from the current observation x_t and the target g without a stochastic latent space. (11) Ours-VallinaGen is a variant of ours, in which the latent space $z \sim q(z|x_t, q)$ is constrained by the standard normal distribution prior p(z).

³https://github.com/pathak22/zeroshot-imitation

We train and evaluate these models on the datasets described in Section III-A. Except for TD-A3C and SAVN, all alternatives are trained with supervision from shortest paths of navigation tasks, although they use different methods, e.g., behavioral cloning as TD-A3C(BC), setting waypoints from experts as GSP, predicting navigation action cost as TD-Semantic, generating next expected observation as NeoNav, etc. We incorporate supervision into RL frameworks for mapless visual navigation, since RL for high-dimensional observations empirically suffers from sample inefficiency [11]. We evaluate these models on two metrics, success rate (SR) and success weighted by (normalized inverse) path length (SPL), as defined in [17]. *More details are provided in [31]*.

TABLE I
Average navigation performance (SR and SPL in $\%$)
COMPARISONS ON UNSEEN SCENES FROM AI2-THOR.

Evaluations	Models	All		$L \ge 5$	
		SR	SPL	SR	SPL
	Random	1.2	0.7	0.6	0.3
	TD-A3C [1]	20.0	4.0	12.9	2.6
	TD-A3C [1](BC)	23.0	7.9	13.4	3.7
Unseen,	Gated-LSTM-A3C [4](BC)	29.1	10.5	19.2	5.1
scenes	GSP [30]	34.4	12.5	27.9	8.3
Known	NeoNav [6]	30.2	11.9	23.6	10.1
targets	Ours	45.7	25.8	41.9	24.8
P=17.7%	Ours-NoVal	34.9	18.1	27.3	14.5
	Ours-FroView	32.3	10.3	29.8	9.4
	Ours-NoGen	41.2	23.8	38.5	22.2
	Ours-VallinaGen	37.5	17.7	34.0	15.9
	Random	2.0	1.0	0.6	0.4
	TD-A3C [1]	10.1	1.9	6.3	1.1
	TD-A3C [1](BC)	12.3	2.4	7.5	1.6
Unseen	Gated-LSTM-A3C [4](BC)	30.0	11.4	26.7	8.6
scenes,	GSP [30]	27.5	8.3	23.4	6.7
Novel	NeoNav [6]	27.4	13.1	22.1	9.3
targets	Ours	37.7	20.5	35.4	19.7
P=16.0%	Ours-NoVal	31.6	10.3	28.9	9.4
	Ours-FroView	24.6	7.8	23.0	6.9
	Ours-NoGen	35.7	19.1	31.6	17.4
	Ours-VallinaGen	31.4	13.9	29.4	12.7

B. Results on the AI2-THOR

Generalization. We analyze the cross-target and crossscene generalization ability of these models on AI2-THOR. The evaluation is divided into two different levels on our testing set, {Unseen scenes, Known targets} and {Unseen scenes, Novel targets}. Each level of evaluation contains 1000 different navigation tasks. [5] proposes using the ratio of the shortest path distance to the Euclidean distance between the start and goal positions to benchmark navigation task difficulty. In each evaluation, we compute the percentage P of the tasks that have a ratio within the range of [1, 1.1] and evaluate the performance on all tasks and on tasks where the optimal path length L is at least 5.

Table I summarizes the results. First, we observe a higher generalization performance for the model with supervision by comparing the results from TD-A3C and TD-A3C(BC). We believe that it is more challenging for RL networks to discover the optimal outputs in the higher-order control tasks. Subsequently, considering the navigation performance difference between TD-A3C(BC) and ours, we see that the idea of generating the future before acting and acting based on the visual difference, works better than directly learning a mapping from raw images to a navigation action. We also compare ours with Gated-LSTM-A3C(BC), which uses an LSTM-based memory and has access to shortest paths during training

as in our method. Our model can consistently outperform the LSTM-based baseline. The model, GSP, also trained with an inverse dynamics model, acquires relatively better navigation performance compared to TD-A3C(BC) and Gated-LSTM-A3C(BC). However, our model gets better results compared with the baseline, which indicates the proposed information-theoretic regularization brings us better generalization for unseen scenes and novel objects. In addition, our method shows at least 10.3% improvement in average SR and 7.4% improvement in average SPL over NeoNav for both cross-target and cross-scene evaluation. The results demonstrate that our regularized navigation model possesses more advantages in target-driven visual navigation.

To further compare ours with SAVN [2], we adapt our navigation model to be driven by a target object class, which is given as a Glove embedding, as in SAVN. The experiment is conducted on AI2-THOR with the same training/testing split, evaluation navigation tasks, and success criterion as [2]. Our model outperforms SAVN with navigation performance (SR and SPL in %) increased by 3.9/7.1 for all navigation trajectories (45.8/23.8 vs. 41.9/16.7) and 13.8/8.2 for trajectories of at least length 5 (43.1/22.4 vs. 29.3/14.2). Our information-theoretic regularization supports both more effective and more efficient navigation.



Fig. 3. We report SR and SPL performance as a function of starting geodesic distance from the target.

Ablation. The ablation on different inputs (Front-view vs. Multi-view) demonstrates that it is easier to generate the next observation when the current information is rich. We also conduct the ablation with four history frames as current inputs, which is difficult to converge during training. We consider that there is no direct connection between the random history and the next observation, which is most related to the current observation and the target. Hence, it is reasonable to generate the next observation from the current multi-view observation rather than from the history frames.

It is difficult for a navigation agent to learn to issue a stop action at a correct location, since there is only one situation with stop action but many cases with other actions during a navigation task, leading to training data imbalance. We first show the cross-target and cross-scene navigation results in a simpler case where the stop signal is provided by the environment rather than being predicted by an agent. The navigation performances (SR and SPL in %) of our model and Ours-NoVal are: 41.6/27.2 and 38.7/19.4, respectively. As expected, both models demonstrate higher performances than the performances in Table I (37.7/20.5 and 31.6/10.3, respectively). Additionally, we see the performance gap of our model is much smaller than Ours-NoVal, meaning we handle the data imbalance better. We consider that the value prediction is critical in learning to issue a stop action. The different stages in navigation tasks can be distinguished by their discounted accumulative reward in RL, and the stage close to the target with a large accumulative reward updates the policy more, which eases the data imbalance.

Based on the ablation on the generation process, we conclude that learning a stochastic latent space is often more generalizable than learning a deterministic one (Ours vs. Ours-NoGen), since the former explicitly models the uncertainty over visual images. However, when the latent space is overregularized by the standard normal distribution prior, the situation is worse (Ours-VallinaGen vs. Ours).

Geodesic distance. We further analyze the navigation performance as a function of the geodesic distance between the start and the target locations in Figure 3. This is based on the 1000 tasks from the {*Unseen scenes*, *Known targets*} evaluation. As can be seen, the geodesic distance is highly correlated with the difficulty of navigation tasks and the performance of all methods degrades as the distance between the start and the target increases. Our model outperforms all alternatives in most cases.



Fig. 4. Average reward curves (left) on AVD during training with a zoomedin view (right). Our model outperforms these alternative navigation methods in terms of learning speed and the final average reward.

C. Results on the AVD

To evaluate the generalization ability in the real world, we train and evaluate our model and some alternatives based on the training and testing splits on AVD.

Navigation driven by target images. We first present the navigation results driven by target images. Figure 4 shows the average returns during training from TD-A3C, TD-A3C(BC), Gated-LSTM-A3C(BC), GSP, and our model. We train five different instances of each algorithm with different random seeds, with each performing one evaluation every 200 navigation episodes. We plot the average return curves with error bands representing the standard deviation. The results show that navigation models with supervision (e.g., TD-A3C(BC), Gated-LSTM-A3C(BC), GSP) learn considerably faster than a pure RL-based navigation model (e.g., TD-A3C). Our proposed model outperforms these baselines both in terms of sample efficiency and the final performance. In Table II, we also report the average values of success rate and SPL with standard deviations of 1000 navigation tasks (P = 15.0%) randomly sampled from unseen scenes in AVD. We observe that all five learning models demonstrate average performance decreases compared to the results on AI2-THOR in Table I,

since both the training scenes and the target views in AVD are limited and the real-world scenes are more complex in structure than synthetic scenes. In addition, our model has relatively larger variance than TD-A3C(BC) and GSP while retaining a better navigation performance than other listed baselines. We consider that our intermediate generative process increases the stochasticity of navigation control, but the proposed information-theoretic regularization generally brings more useful information for navigation tasks, which is critical for policy learning from the perceptible environment.

We visualize navigation tasks from our model in [31]. These tasks are all characterized by unknown scenes and long distances between the start points and the targets. Our agent can navigate to the targets successfully in the first three tasks. In the last three tasks, our model fails and the problems include thrashing around in space without making progress, getting stuck in the corridor, and navigating around tight spaces (e.g, the small bathroom where the third task starts).

TABLE II AVERAGE NAVIGATION PERFORMANCE (SR, SPL and CR in %) comparisons on unseen scenes from AVD and realistic

SCENARIOS (KS).								
	AV	RS						
Model	SR	SPL	SR	CR				
Random	$2.8_{(0.9)}$	$1.8_{(0.4)}$	2.0	62.0				
TD-A3C [1]	9.3(2.4)	2.9(1.1)	2.0	58.0				
TD-A3C [1](BC)	$15.9_{(1.9)}$	6.1(0.9)	8.0	56.0				
Gated-LSTM-A3C [4](BC)	$13.3_{(2.7)}$	$5.8_{(1.6)}$	10.0	46.0				
GSP [30]	19.3(1.1)	$5.5_{(0.6)}$	24.0	48.0				
Ours	23.1 _(2.1)	13.5 (1.0)	28.0	40.0				

Navigation driven by target labels. We also adapt our method to compare it with TD-Semantic [10], in which the navigation goal is defined in the form of a one-hot vector over a prescribed set of semantic labels, {Couch, Table, Refrigerator, Microwave, TV}. The experiment is conducted on AVD with the same training/testing split, evaluation tasks, and success criterion as [10]. While sharing the same idea of improving the training by using the supervision from shortest paths of navigation tasks, our method outperforms TD-Semantic by 22% for RGB input $(53\% \ vs. \ 31\%)$, and 28% for depth input $(59\% \ vs. \ 31\%)$ for average success rate on the AVD testing set. Our method (with depth input) shows a 5.4% improvement in average success rate compared to TD-Semantic with semantic input (59.0% vs. 53.6%), which is provided by some state-ofthe-art detectors and segmentors. The best performances of two methods over various target labels are presented in [31]. We suggest that our information-theoretic regularization helps learn a controller that can analyze the relationship between visual observation and the target and extract some important information to guide navigation. This process is not affected by the target format, e.g., a semantic label or a view image.

D. Results on the Real World

Moving to the real-world scenarios further shows the generalization capabilities of the proposed navigation models and the robustness against indoor layouts, robot types and sensor types. The models are trained based purely on the discrete dataset (e.g., AVD), and the real-world environments are continuous and unknown to the agents. **Robotic setup.** We demonstrate the proposed model using a TurtleBot. The configuration of the TurtleBot is shown in Figure 5(a), which consists of a differential wheeled moving base Kobuki and four RGB Monocular cameras equipped at the top of the robot. The proposed system takes as input data from four real-time camera sensors and a target image at each time step, to predict the optimal navigation action. The action command is converted to the wheel velocity and passed to the robot. For example, the *move right* action in \mathcal{A} is converted to rotate right at $45^{\circ}/s$ for 2s, move forward at 0.25m/s for 2s, and rotate left at $45^{\circ}/s$ for 2s. These commands are published with a frequency of 5Hz. It is complex due to the movement direction restrictions of the TurtleBot.

Transfer to the real world. Experiments were conducted on a floor (approx. $400m^2$) of an academic building. We evaluate the robot with 50 randomly sampled navigation tasks in the scene. A navigation task is regarded as a success only if the robot stops near the target (e.g., 1m) within 100 steps, and we consider it a failure if the robot collides with an obstacle or does not reach the goal within the step limit. Although the model is trained on the discrete household dataset, it can transfer to the realistic public scenarios and exhibit robustness towards random starting points, varying step lengths, changes in illumination, target types, scene layouts, etc. A quantitative analysis of these navigation tasks is provided in Table II (the right-most column), where the average success rates and collision rates (CR, the rate of collision cases to all navigation cases) are listed. We observe that the existence of the proposed information-theoretic regularization enables better transfer to new environments. However, all models present consistently high collision rates during navigation, since realistic evaluation, characterized by continuous space and robotic movement deviation, is very challenging. Extension to depth input or simultaneous mapping [33] would make the method applicable in more general scenarios. Figure 5 qualitatively compares our method with the GSP baseline. The baseline generally gets stuck behind the obstacle (e.g. the wall) and tries hard to move forward, while our method finds the way towards the door and issues the stop action close to the target.

V. CONCLUSION

We propose integrating an information-theoretic regularization into a deep reinforcement learning framework for the target-driven task of visual navigation. The regularization maximizes the mutual information between navigation actions and visual observations, which essentially models the actionobservation dynamics and the causality between navigation actions and observational changes. By adapting the regularization for target-driven navigation, our agent further learns to build the correlation between the observation and the target. The experiments on the simulation and the real-world dataset show that our model outperforms some baselines by a large margin in both the cross-scene and the cross-target navigation generalization. Experiments using the TurtleBot robot demonstrate the transfer capability of the proposed navigation model, which is easy to deploy on-robot.

In this work, training requires supervision from expert trajectories, which are generated based on the topology graphs of training scenes. For future work, we will investigate how real-world human demonstrations can be leveraged and how the model can be extended to dynamic environments [34]. In addition, extending some state-of-the-art model-free and model-based deep RL algorithms to the target-driven visual navigation problem for exploring better navigation generalization is also a great topic for future work.



Fig. 5. Qualitative examples. (a) The robotics system setup and the navigation task. (b) The trajectory based on GSP. (c) The trajectory based on our model.

References

- Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *Proc. ICRA*, 2017, pp. 3357–3364.
- [2] M. Wortsman, K. Ehsani, M. Rastegari, A. Farhadi, and R. Mottaghi, "Learning to learn how to learn: Self-adaptive visual navigation using meta-learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 6750–6759.
- [3] W. Sun, A. Venkatraman, G. J. Gordon, B. Boots, and J. A. Bagnell, "Deeply aggrevated: Differentiable imitation learning for sequential prediction," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 3309–3318.
- [4] Y. Wu, Y. Wu, G. Gkioxari, and Y. Tian, "Building generalizable agents with a realistic and rich 3d environment," in 6th International Conference on Learning Representations, ICLR 2018, Workshop Track Proceedings. OpenReview.net, 2018.
- [5] M. Savva, J. Malik, D. Parikh, D. Batra, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, and V. Koltun, "Habitat: A platform for embodied AI research," in 2019 IEEE International Conference on Computer Vision. IEEE, 2019, pp. 9338–9346.
- [6] Q. Wu, D. Manocha, J. Wang, and K. Xu, "Neonav: Improving the generalization of visual navigation via generating next expected observations," in *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI* 2020. AAAI Press, 2020, pp. 10001–10008.
- [7] S. Bai, J. Wang, F. Chen, and B. Englot, "Information-theoretic exploration with bayesian optimization," in 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2016, pp. 1816–1822.
- [8] W. Yang, X. Li, and Z. Zhang, "A regularized approach to sparse optimal policy in reinforcement learning," in *Advances in Neural Information Processing Systems*, 2019, pp. 5938–5948.
- [9] S. Mohamed and D. J. Rezende, "Variational information maximisation for intrinsically motivated reinforcement learning," in *Advances in neural information processing systems*, 2015, pp. 2125–2133.
- [10] A. Mousavian, A. Toshev, M. Fišer, J. Košecká, A. Wahid, and J. Davidson, "Visual representations for semantic target driven navigation," in 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019, pp. 8846–8852.
- [11] L. Kaiser, M. Babaeizadeh, P. Milos, and et al., "Model based reinforcement learning for atari," in 8th International Conference on Learning Representations, ICLR 2020. OpenReview.net, 2020.
- [12] J. Liang, U. Patel, A. J. Sathyamoorthy, and D. Manocha, "Crowd-steer: Realtime smooth and collision-free robot navigation in densely crowded scenarios trained using high-fidelity simulation," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, *IJCAI 2020*, 2020, pp. 4221–4228.

- [13] A. J. Sathyamoorthy, J. Liang, U. Patel, T. Guan, R. Chandra, and D. Manocha, "Densecavoid: Real-time navigation in dense crowds using anticipatory behaviors," in 2020 IEEE International Conference on Robotics and Automation (ICRA), 2020, pp. 11345–11352.
- [14] F. Xia, W. B. Shen, C. Li, P. Kasimbeg, M. E. Tchapmi, A. Toshev, R. Martín-Martín, and S. Savarese, "Interactive gibson benchmark: A benchmark for interactive navigation in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 713–720, 2020.
- [15] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," in *Proc. CVPR*, 2017, pp. 2616–2625.
- [16] N. Savinov, A. Dosovitskiy, and V. Koltun, "Semi-parametric topological memory for navigation," in 6th International Conference on Learning Representations, ICLR 2018, Conference Track Proceedings. OpenReview.net, 2018.
- [17] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Mottaghi, "Visual semantic navigation using scene priors," in *7th International Conference* on Learning Representations, ICLR 2019. OpenReview.net, 2019.
- [18] N. Trigoni, A. Markham, and L. Xie, "Snapnav: learning mapless visual navigation with sparse directional guidance and visual reference," in *IEEE International Conference on Robotics and Automaton (ICRA)*. IEEE Digital Library.
- [19] N. Hirose, F. Xia, R. Martín-Martín, A. Sadeghian, and S. Savarese, "Deep visual mpc-policy learning for navigation," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3184–3191, 2019.
- [20] Y. Gao, H. Xu, J. Lin, F. Yu, S. Levine, and T. Darrell, "Reinforcement learning from imperfect demonstrations," in 6th International Conference on Learning Representations, ICLR 2018, Workshop Track Proceedings. OpenReview.net, 2018.
- [21] J. Ho and S. Ermon, "Generative adversarial imitation learning," in Advances in neural information processing systems, 2016, pp. 4565– 4573.
- [22] G. Li, M. Müller, V. Casser, N. Smith, D. L. Michels, and B. Ghanem, "OIL: observational imitation learning," in *Robotics: Science and Systems XV, 2019*, A. Bicchi, H. Kress-Gazit, and S. Hutchinson, Eds., 2019.
- [23] M. Muller, G. Li, V. Casser, N. Smith, D. L. Michels, and B. Ghanem, "Learning a controller fusion network by online trajectory filtering for vision-based uav racing," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.
- [24] S. MartíNez and F. Bullo, "Optimal sensor placement and motion coordination for target tracking," *Automatica*, vol. 42, no. 4, pp. 661– 668, 2006.
- [25] A. Dame and E. Marchand, "Mutual information-based visual servoing," *IEEE Transactions on Robotics*, vol. 27, no. 5, pp. 958–969, 2011.
- [26] B. Charrow, G. Kahn, S. Patil, S. Liu, K. Goldberg, P. Abbeel, N. Michael, and V. Kumar, "Information-theoretic planning with trajectory optimization for dense 3d mapping." in *Robotics: Science and Systems*, vol. 11, 2015.
- [27] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 16–17.
- [28] S. Racanière, T. Weber, D. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. Heess, Y. Li *et al.*, "Imaginationaugmented agents for deep reinforcement learning," in *Advances in neural information processing systems*, 2017, pp. 5690–5701.
- [29] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution," in Advances in Neural Information Processing Systems, 2018, pp. 2450–2462.
- [30] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell, "Zero-shot visual imitation," in *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition Workshops, 2018, pp. 2050–2053.
- [31] Q. Wu, K. Xu, J. Wang, M. Xu, and D. Manocha, "Reinforcement learning based visual navigation with information-theoretic regularization," 2019. [Online]. Available: http://arxiv.org/abs/1912.04078
- [32] J. Zhang, J. T. Springenberg, J. Boedecker, and W. Burgard, "Deep reinforcement learning with successor features for navigation across similar environments," in 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2017, pp. 2371–2378.
- [33] D. S. Chaplot, D. Gandhi, A. Gupta, and R. R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," in Advances in Neural Information Processing Systems 33, NeurIPS 2020, 2020.
- [34] A. J. Sathyamoorthy, U. Patel, T. Guan, and D. Manocha, "Frozone: Freezing-free, pedestrian-friendly navigation in human crowds," 2020. [Online]. Available: https://arxiv.org/abs/2003.05395