

NeoNav: Improving the Generalization of Visual Navigation via Generating Next Expected Observations

Qiaoyun Wu,^{1,2} Dinesh Manocha,² Jun Wang,¹ Kai Xu^{3*}

¹Nanjing University of Aeronautics and Astronautics, ²The University of Maryland

³National University of Defense Technology

(Project page: <http://kevinkaixu.net/projects/neonav.html>)

Abstract

We propose improving the cross-target and cross-scene generalization of visual navigation through learning an agent that is guided by conceiving the next observations it expects to see. This is achieved by learning a variational Bayesian model, called *NeoNav*, which generates the *next expected observations (NEO)* conditioned on the current observations of the agent and the target view. Our generative model is learned through optimizing a variational objective encompassing two key designs. First, the latent distribution is conditioned on current observations and the target view, leading to a model-based, target-driven navigation. Second, the latent space is modeled with a Mixture of Gaussians conditioned on the current observation and the next best action. Our use of mixture-of-posteriors prior effectively alleviates the issue of over-regularized latent space, thus significantly boosting the model generalization for new targets and in novel scenes. Moreover, the NEO generation models the forward dynamics of agent-environment interaction, which improves the quality of approximate inference and hence benefits data efficiency. We have conducted extensive evaluations on both real-world and synthetic benchmarks, and show that our model consistently outperforms the state-of-the-art models in terms of success rate, data efficiency, and generalization.

1 Introduction

Mapless visual navigation is an important skill for robots operating in unknown, unstructured environments. It is characterized as the ability of a robot to navigate itself from an arbitrary location in the environment to a goal position, based solely on the visual inputs from its on-board sensors. The main challenge of visual navigation lies in understanding the scene layout based on the visual observations and reasoning about the spatial relation between the current and the target location. This is the main impediment hindering the generalization of navigation ability across different scenes due to visual and structural discrepancy.

Recent years have witnessed fast advancement of visual navigation thanks to deep learning, e.g., deep reinforcement learning (RL) models (Zhang et al. 2017; Zhu et al. 2017; Gupta et al. 2017; Racanière et al. 2017). Model-free approaches learn to directly map raw observations to values or actions, which usually suffers from low data efficiency.

Model-based methods tackle this issue through modeling the transition dynamics of agent-environment interaction. Such model can be used to reason about the future, thus relieving the trial-and-error learning endeavor. However, it is difficult to learn a powerful model that generalizes across different scenes, which is known as the model imperfection issue.

We propose *NeoNav*, a model-based, supervised learning approach to visual navigation with strong model generality. In our method, the agent is guided by conceiving the next observations it expects to see supposing the best action is taken. This is realized by learning a *generative model* conditioned on the multi-view observations at the current location as well as the target view, from which the next expected observation (NEO) can be generated. We predict the next best action based on the generated NEO and the current (front-view) observation. We frame this problem as a variational Bayesian inference where the variational lower-bound (objective) consists of three terms: reconstruction, regularization and classification. The minimization of the reconstruction error maximizes the likelihood of the NEO given the current observations and target view. The regularization term drives the variational posterior to match a prior distribution. The classification term is devised for action prediction.

The key characteristic of our approach is the modeling of the latent space. First, to enable target-driven navigation, the latent distribution is conditioned on current observations and target views. Second, we model the latent space with a Mixture of Gaussians conditioned on current observations and next best actions. Such a variational *mixture of posteriors* prior (Tomczak and Welling 2017) effectively alleviates over-regularization of the latent space, thus facilitating cross-scene model generalization. Moreover, NEO generation via sampling over the latent space essentially models the *forward dynamics* of the agent-environment interaction, i.e. action-driven state transition. This improves the expressiveness of the approximation of variational inference (Cremer, Li, and Duvenaud 2018), thereby greatly enhancing inference generalization and data efficiency.

Although our supervised setting requires denser training signals, the learned model shows significantly better generality over unseen test scenes than RL-based approaches (even if enhanced by supervision such as in behavior cloning). Fortunately, target-driven navigation tasks enjoy the easy acquisition of ground-truth paths as training data (e.g., using A* algorithm). We conducted extensive evaluations on public datasets of both synthetic (AI2-THOR

*Corresponding author: Kai Xu (kevin.kai.xu@gmail.com)
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

framework (Zhu et al. 2017)) and real-world (Active Vision Dataset, AVD (Ammirato et al. 2017)) scenes. We demonstrate that our model attains at least a 5% higher success rate for both cross-target and cross-scene generalization, compared to several state-of-the-art alternative methods, ranging from model-based to model-free, from RL-based to supervised, and from target-driven to semantic-driven.

2 Related Works

Model-free navigation. This approach learns to map the raw observations directly to actions. Mnih et al. (2015) present the first deep reinforcement learning model, called Deep Q-Learning, that successfully learns control policies directly from high-dimensional sensory input. Schulman et al. (2015) propose the Trust Region Policy Optimization (TRPO), which is effective for optimizing large nonlinear policies and demonstrates robust performance on a wide variety of robotic tasks. Lillicrap et al. (2015) present the DDPG (Deep Deterministic Policy Gradient), which can robustly solve many simulated tasks. Model-free methods usually require large training data and the policies do not readily generalize to novel tasks in unseen environments.

Several works study using deep neural networks to realize classical iterative planning algorithms without an explicit environmental model (Tamar et al. 2016; Silver et al. 2017; Oh, Singh, and Lee 2017; Lee et al. 2018). Zhang et al. (2017) focus on the problem of robot navigation in maze-like environments and present a successor-feature-based deep RL algorithm that can transfer navigation policies across similar environments. Most existing models are trained and tested on mazes; an exception is (Mirowski et al. 2018) which proposes a DRL model for navigating in cities.

Zhu et al. (2017) propose an excellent feed-forward architecture for target-driven visual navigation by combining a Siamese network with the A3C algorithm (Mnih et al. 2016). They focus on cross-target generalization in smaller indoor scenes and do not consider generalization to previously unseen environments. In (Mousavian et al. 2018), semantic scene segmentation is incorporated in learning to map from semantic information to navigation actions. Through comparison, we show that our method has better cross-target and cross-scene generalization.

Model-based navigation. This approach achieves better data efficiency, but has the issues of cross-scene generalization due to model imperfections. Several approaches have been proposed to address the model imperfection issue, such as capturing model uncertainty (Deisenroth and Rasmussen 2011; Marco et al. 2017) and incorporating semantic priors into environmental models (Yang et al. 2018; Mousavian et al. 2018). With the advances of attention mechanisms in deep learning, many works propose modeling the environment with a memory unit.

Savinov et al. (2018) introduce a topological landmark-based memory for navigation. A common issue with such an approach is that the memory, representing an allocentric map of the scene, grows in size as the scene exploration proceeds, limiting its practical utility in navigating within

large environments. In (Gupta et al. 2017), the problem is alleviated by learning an ego-centric mapper and planner, which, however, assumes perfect odometry. Henriques and Vedaldi (2018) develop a differentiable module that is able to associate an egocentric representation of a scene to an allocentric one. Our method models the environment with the probabilistic latent distribution in a variational Bayesian framework, where both model generality and model scalability are attained by imposing a mixture-of-posteriors prior. A similar model was proposed in (Henriques and Vedaldi 2018) in the RL setting.

Imagination-based navigation. Some model-based navigation methods reason about the future based on the internal model. Razvan et al. (2017) introduce an imagination-based planner, which is the first model-based, sequential decision-making agent that can learn to propose (imagine), evaluate, and execute plans. The method demonstrates good performance on 2D maze-solving tasks. Imagination-Augmented Agents (I2As) (Racanière et al. 2017) was later proposed and it learns to generate and interpret predictions as additional context for deep policy networks. These methods are generally data-efficient, but have difficulty in scaling to complex, high-dimensional tasks. Generally, similar ideas have been well practiced in the studies of deep RL, where the internal model is used to predict future observations and/or rewards (Oh et al. 2015; Leibfried, Kushman, and Hofmann 2016; Dosovitskiy and Koltun 2016; Jaderberg et al. 2016; Mirowski et al. 2016; Finn and Levine 2017; Pathak et al. 2017). Watter et al. (2015) introduce Embed to Control (E2C), which learns to generate image trajectories from a latent space in which the dynamics is constrained to be locally linear, in contrast to the non-linear dynamics modeled by our latent space. Although sharing a similar spirit, our work is significantly different from the imagination-augmented RL-based navigation (Racanière et al. 2017; Pascanu et al. 2017). First, our model is formulated as a variational Bayesian inference trained with supervised learning rather than RL. Second, their imaginations refer to a simulated rollout of trajectories, while our imagination is a one-step imagination of the next observation.

3 Method

3.1 Problem Setting

Target-driven visual navigation takes the current observations x captured by the agent and a target view g as input, and predicts the next best action a at each time step to navigate the agent, until reaching the target position.

Observations and goals. The agent camera has only the azimuth DoF. At each agent location, the observation x consists of K views with evenly distributed azimuth angles: $\{0^\circ, \frac{1}{K}360^\circ, \dots, \frac{K-1}{K}360^\circ\}$, in which 0° corresponds to the front-looking view. The agent captures an image (RGB, depth or RGB-D) at each view. The K -view observations provide a local context of the environment, based on which the agent is able to reason about its location and the room

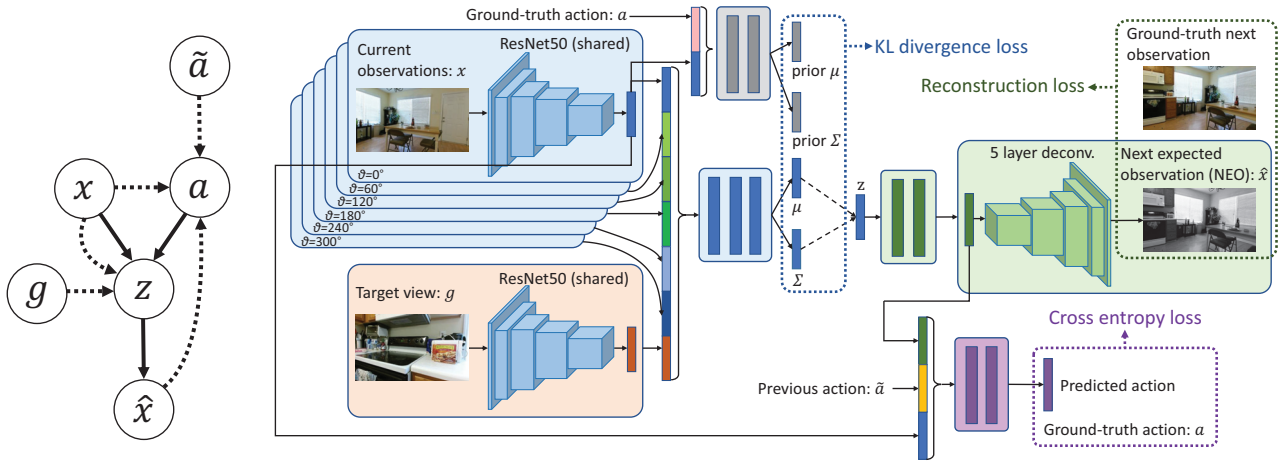


Figure 1: Model overview: the probabilistic graphical model and the network architecture. In the graphical model, the posterior $p_\theta(z|x, a)$ and generator $p_\theta(\hat{x}|z)$ are denoted with solid lines, while the variational approximation $q_\lambda(z|x, g)$ and the action prediction $q_\varphi(a|x, \hat{x}, \tilde{a})$ are depicted with dashed lines. The generative model is realized with a variational auto-encoder architecture: The encoder takes the current observations x and the target view g as input. The decoder generates the NEO \hat{x} from a random vector sampled from the latent space defined by the Gaussian $\mathcal{N}(\mu, \Sigma)$. The feature in the decoding module is used in predicting the next action. Three losses used for learning the generative model are marked with dashed boxes.

layout of its surroundings. The target view is consistent with the observation views in terms of image data modality.

Action space. At each time step, the agent can choose one action from a discrete set of allowable actions: $\{\text{move_forward}, \text{move_back}, \text{move_left}, \text{move_right}, \text{rotate_ccw}, \text{rotate_cw}, \text{stop}\}$, where *move* means horizontal movement of the agent and *rotate* refers to azimuth rotation of the camera. *ccw* and *cw* stands for counter-clockwise and clockwise, respectively. If there is no next observation view associated with an action, the action is considered to cause a collision.

3.2 The Variational Bayesian Navigation Model

Given the current observation x , instead of directly predicting the next best action a as in many other works, we opt to first *generate* the next expected observation (NEO) \hat{x} assuming that the next best action a is known *a priori* and is executed. This can be described with a generative model:

$$p_\theta(\hat{x}, z|x, a) = p_\theta(\hat{x}|z)p_\theta(z|x, a), \quad (1)$$

where $p_\theta(\hat{x}, z|x, a)$ is a parametric model of the joint distribution over the NEO \hat{x} and a latent variable z . Essentially, this generative model is a probabilistic *forward dynamics model* of the agent, where the acquirement of the next observation is driven by the selected next action.

To learn the generative model, one typically maximizes the marginal log-likelihood $\log p_\theta(\hat{x}|x, a)$. However, when the model is parameterized by a neural network, the optimization could be difficult due to the intractability of the marginal likelihood. Moreover, the next best action a is unknown *a priori* and is inherently determined by the target g . To this end, we apply *variational inference* and introduce an

inference network $q_\lambda(z|x, g)$ with parameters λ to approximate the true posterior $p_\theta(z|x, a)$. In particular, we optimize the following lower bound of the marginal likelihood:

$$\log p_\theta(\hat{x}|x, a) \geq \mathbb{E}_{z \sim q_\lambda(z|x, g)} [\log \frac{p_\theta(\hat{x}, z|x, a)}{q_\lambda(z|x, g)}] = \mathcal{L}(\hat{x}). \quad (2)$$

This lower bound forms our objective function:

$$\mathcal{J} = - \mathbb{E}_{z \sim q_\lambda(z|x, g)} [\log p_\theta(\hat{x}|z)] + KL[q_\lambda(z|x, g) || p_\theta(z|x, a)] = -\mathcal{L}(\hat{x}), \quad (3)$$

where KL denotes the Kullback-Leibler divergence. During training, $p_\theta(z|x, a)$ can be estimated as a Gaussian distribution conditioned on the current observation x and the ground-truth action a , leading to a *mixture-of-posteriors prior* imposed on the latent distribution.

To realize robot navigation, we learn a navigation action classifier $q_\varphi(a|x, \hat{x}, \tilde{a})$ which predicts the next best action a based on the current observation x , the generated NEO \hat{x} as well as the previous action \tilde{a} . Integrating action prediction, the objective function becomes:

$$\mathcal{J} = - \alpha \mathbb{E}_{z \sim q_\lambda(z|x, g)} [\log p_\theta(\hat{x}|z)] + \beta KL[q_\lambda(z|x, g) || p_\theta(z|x, a)] + \gamma \mathbb{E}_{a \sim p(a)} [-\log q_\varphi(a|x, \hat{x}, \tilde{a})], \quad (4)$$

where $a \sim \text{Cat}(1/C)$. A complete derivation of this objective is given in the supplemental material. The objective function in (4) is composed of a reconstruction loss, a KL divergence loss and a cross entropy loss. The three hyperparameters are empirically set as $\alpha = 0.01$, $\beta = 0.0001$ and $\gamma = 1$ throughout our experiments. Figure 1(left) shows the probabilistic graphical model of our navigation model.

3.3 The Network Architecture

Corresponding to the variational objective, the architecture of our network consists of three subnetworks (see Figure 1). The *variational inference module* takes the full observation views at the current robot position as well as the target view as input and extracts a 2048-D feature vector for each of them using a ResNet-50. The input image resolution is $64 * 64$. These output 2048-D feature vectors are then used to infer a vector of latent variables of dimension 400 with a MLP. Here, a KL divergence loss is minimized to impose the distribution of the latent variables to match a prior estimated from the current observation (front view only) and the ground-truth action. The *NEO generation module* then generates the NEO in the front view out of a latent vector, using a two-layer MLP followed by a 5-layer convolutional network (please refer to the supplemental material for details). This task is trained with the supervision of ground-truth next observation. The *action prediction module* maps the concatenation of the last layer feature of the NEO generation module (2048-D), the feature of the current observation (2048-D) and the feature (1024-D) extracted from the previous action (7-D one-hot vector) into the predicted next action (7-D), using a four-layer MLP. Ground-truth actions are used to train this subnetwork.

Model training and testing. Our model is trained and tested with both real-world environments from the Active Vision Dataset (AVD) (Ammirato et al. 2017) and synthetic scenes of AI2-THOR (Zhu et al. 2017). Each scene in the dataset is represented as a grid of robot locations (see Figure ??). The size of the grid cell is 0.25-0.5 meters. For each grid point, 6 azimuth camera views are captured for AVD and 4 for AI2-THOR. For the task of target-driven navigation, the ground-truth navigation path is simply the shortest path over the grid. The optimization of the variational objective is achieved by Monte Carlo sampling, where the gradients are backpropagated with the standard reparameterization trick (Kingma and Welling 2013).

At test time, our model is used as a controller for the agent to predict the next action given the current observations. We feed the current observation views and the target view into the inference module to obtain a Gaussian component in the latent space. The features used for NEO prediction and extracted for the current front-view observation, as well as the previous action, are used for next action prediction. The actual generation of NEO, however, is not needed in testing.

4 Experiments

We evaluate both cross-target and cross-scene generalization, as well as a few other important characteristics, of our model by comparing it with one baseline and a few state-of-the-art methods. We also compare our method to two ablated variants of it to justify our major design choices. In addition, we visualize the latent space of our model for a better understanding of what we learn, as well as the navigation paths for a qualitative evaluation.

Experimental settings. Our evaluations are conducted on both AVD and AI2-THOR. AVD contains 11 relatively complex real-world houses, of which 8 houses were used for training and 3 for testing. AI2-THOR contains 120 scenes in four categories including kitchen, living room, bedroom, and bathroom. Each category includes 30 scenes, out of which 20 are used for training and 10 for testing. For all the methods being compared, we train a single model for all the AVD scenes and separate models for the categories of AI2-THOR.

For each training scene, we choose fifteen different views as the target, each of which contains a targeted object such as a dining table, a refrigerator, a sofa, a television, etc. During testing, the target views are randomly sampled from the test scenes, encompassing both views similar to trained targets and views unseen in training. When sampling start points, we consider the ratio of the shortest path distance to the Euclidean distance between start and goal positions (Savva et al. 2019). We perform aggressive rejection sampling to ensure that 15% of the tasks have a ratio within the range of $[1, 1.1]$. As in many navigation systems, a collision detection module is devised. When a collision is detected, the action with the next largest probability is chosen.

Success criteria. In each episode, the agent runs until arriving at the goal (the distance to the target position is less than 1 meter and the angle between the current and target view direction is less than 90°), reaching the maximum number of steps (100), or issuing a *stop* action. In the setting with *stop* action, an episode is successful if and only if the agent issues a *stop* action exactly when it reaches the goal. This success criterion is apparently stricter than that in the setting without a *stop* action. We will evaluate both cases.

Evaluation metrics. We adopt two evaluation metrics, success rate and success weighted by (normalized inverse) path length (SPL) (Anderson et al. 2018). Success rate is the fraction of the runs that successfully navigate to the goal. SPL is defined as $\frac{1}{N} \sum_{i=1}^N S_i \frac{L_i}{\max\{P_i, L_i\}}$, where N is the number of navigation tasks, S_i a binary indicator of success in the i -th task. P_i and L_i denote the actual path length and the shortest path distance for the i -th task, respectively.

We compare with the following baselines/alternatives:

- **Random Walk**, a baseline where the agent randomly chooses an action at each time step.
- **TD-A3C**, a target-driven visual navigation model based on deep RL (Zhu et al. 2017). The reactive policy is trained with the views of the previous three steps in addition to the current view. In addition, we did not freeze the ResNet-50 when training the model. We compare against two variants of the model, with or without supervision, denoted as TD-A3C and TD-A3C-U, respectively.
- **I2A**, i.e., Imagination-Augmented Agents (Racanière et al. 2017), which is a model-based deep RL model. The original method is developed for 2D maze-solving tasks; we re-implemented it for visual navigation by changing its input to first-person views.

Table 1: Navigation performance (success rate and SPL, in %) comparison on *novel scenes* from AVD with *stop* action.

Model \ Target	Table	Exit	Couch	Refrigerator	Sink	Avg.
Random Walk	4.0 / 2.7	4.6 / 3.1	3.4 / 2.1	3.2 / 2.1	3.6 / 2.7	3.8 / 2.7
TD-A3C-U	5.3±0.6 / 2.4±0.3	6.7±0.4 / 4.2±0.2	4.3±0.4 / 2.5±0.2	5.6±0.3 / 2.7±0.2	7.1±0.5 / 3.6±0.2	5.8 / 3.1
TD-A3C	12.4±2.1 / 1.5±0.6	23.0±1.6 / 2.8±0.5	15.0±1.9 / 1.7±0.2	7.2±1.1 / 1.1±0.2	13.4±1.4 / 1.7±0.3	14.2 / 1.8
Gated-LSTM-A3C	6.5±0.7 / 1.3±0.3	16.2±0.8 / 4.0±0.2	8.0±0.5 / 1.2±0.3	14.3±0.7 / 3.3±0.4	6.5±0.5 / 0.9±0.2	10.3 / 2.1
I2A	10.6±1.1 / 1.3±0.3	21.7±1.3 / 2.2±0.4	14.3±0.9 / 2.1±0.3	8.9±0.7 / 2.3±0.3	11.2±0.9 / 1.6±0.3	13.3 / 1.9
Ours	12.1±1.3 / 3.4±0.8	30.4±1.6 / 8.4±0.9	11.0±1.1 / 2.9±0.9	35.0±1.4 / 11.9±1.0	11.0±0.9 / 2.7±0.4	19.9 / 5.9

Table 2: Navigation performance (success rate and SPL, in %) comparison on *novel scenes* from AVD without *stop* action.

Model \ Target	Table	Exit	Couch	Refrigerator	Sink	Avg.
Random Walk	34.8 / 12.9	29.0 / 11.3	29.8 / 10.8	27.4 / 10.7	23.0 / 10.2	28.8 / 11.2
TD-A3C-U	39.7±4.4 / 13.1±2.6	29.3±4.3 / 10.9±2.1	30.1±4.1 / 9.9±2.0	28.4±3.3 / 10.1±1.7	22.9±3.0 / 9.7±1.9	30.1 / 10.7
TD-A3C	45.8±5.3 / 5.8±1.7	37.6±3.6 / 6.3±1.4	37.2±4.0 / 5.0±1.3	16.8±3.2 / 4.4±1.0	23.4±3.7 / 4.7±1.9	32.2 / 5.2
Gated-LSTM-A3C	31.0±3.3 / 8.6±1.0	31.1±2.7 / 13.7±1.5	25.3±2.3 / 5.8±1.1	31.4±2.9 / 12.9±1.9	23.0±2.4 / 8.3±1.0	28.4 / 9.9
I2A	40.6±4.2 / 6.7±2.1	37.3±3.6 / 6.9±2.0	35.9±3.3 / 6.3±1.7	17.3±3.1 / 5.6±1.2	20.1±3.0 / 6.3±1.4	30.2 / 6.4
Ours	57.6±4.5 / 32.6±2.9	52.4±4.1 / 23.9±2.7	43.4±3.9 / 21.1±2.1	46.4±3.4 / 24.3±1.9	38.8±3.3 / 25.8±1.7	47.7 / 25.5
Ours-FrontView	49.9±4.1 / 24.3±1.6	40.6±3.9 / 13.4± 1.2	38.1±4.0 / 16.1± 1.5	28.7±3.4 / 12.1±1.3	28.1±3.5 / 11.9± 2.1	37.1 / 15.6
Ours-NoGen	43.6±2.1 / 28.6±0.9	34.3±2.3 / 19.4±0.7	31.7±1.9 / 18.4±0.7	37.7±2.0 / 20.1±1.1	26.8±2.6 / 22.0±1.5	34.8 / 21.7
Ours-NoMoP	50.7±5.3 / 28.2±3.6	49.9±4.2 / 22.8±3.1	42.9±4.3 / 20.4±2.9	34.4±3.9 / 19.9±2.0	23.7±3.7 / 17.2±2.1	40.3 / 21.7
Ours (RGB)	48.4±4.9 / 25.7±3.1	32.4±3.1 / 16.7±2.0	28.7±2.7 / 16.8±1.4	37.9±4.5 / 18.8±2.9	38.2±2.3 / 26.4±1.1	37.1 / 20.9
Ours (RGBD)	59.6±4.3 / 36.8±3.0	38.4±4.4 / 18.5±2.9	31.5±3.7 / 17.8±1.6	43.5 ±4.1 / 15.3±2.4	37.8±3.2 / 21.7±2.0	42.2 / 22.0

- **Gated-LSTM-A3C**, an LSTM-based variant of A3C model adapted from (Wu et al. 2018), where we train the model with back-propagation through time over 10 unrolled time steps; the goal is specified as an image.
- **TD-Semantic**, a state-of-the-art target-driven navigation model based on deep supervised learning. The method leverages the semantic and contextual representations obtained by off-the-shelf object detection and segmentation methods (Mousavian et al. 2018).
- **Ours-FrontView**, a variant of ours where the current observation at each time step is only the front view rather than four views.
- **Ours-NoGen**, a non-generative variant of our model where the next expected observation is predicted directly from the current observations and the target view. This is implemented simply by removing the Gaussian sampling and the KL-loss in Eq. (4).
- **Ours-NoMoP**, a baseline variant of our model in which the latent space follows the standard normal distribution prior, instead of the mixture-of-posteriors prior.

Except for the Random Walk and TD-A3C-U, all alternatives are trained with supervision. TD-A3C, I2A and Gated-LSTM-A3C are all first trained via behavioral cloning using ground-truth paths. After pre-training, we update the three policy layers using a shaped reward based on the geodesic distance to the goal, $geo(x, g)$, as described in (Gordon et al. 2019): $r_t = geo(x_{t-1}, g) - geo(x_t, g) + \zeta$, where $\zeta = -0.01$ is a small constant time penalty. More implementation details are provided in the supplemental material. Unless explicitly stated otherwise, all methods take depth images as the input observation. For our model, we also implement

variants taking RGB and/or RGBD images as input.

Cross-target generalization. Over the 8 training scenes of AVD, we evaluate navigation performance for 40 novel targets that are unseen in the training phase. These targets are classified into five intervals of the shortest distance between the test and the nearest trained targets: [2, 3], [4, 5], [6, 7], [11, 13], and [14, 16]. For each interval, we sample 1000 navigation tasks with different starting points. The results on the two metrics (with standard deviation measured from five times training) are reported in Figure 2. Generally, the success rate decreases as the distance between the test and trained targets increases. Our model with default depth input outperforms the state-of-the-art alternatives by > 5% for average success rate and by > 4% for average SPL. An observation is that the success rate is related to the degree of presence of the targeted object in the target views. If the targeted object is completely visible in the target view, the target is more instructive and thus easier to reach.

Cross-scene generalization. To evaluate model generality over unseen scenes, we perform navigation with 15 sampled target views from the testing split of AVD. The targets are classified into five groups according to the object of interest in the target views; see Table 1. Note that object labels were not used for navigation. For each group, we sample 1000 navigation tasks (starting points). Our model achieves > 5% higher average success rate and > 2% higher average SPL than the alternative methods. The standard deviations in the table are measured from training each model for five times. In Table 2, we report the results for the case without a *stop* action. The results of the tasks reversing the start and




Figure 2: Navigation performance (success rate and SPL, in %) comparison for novel targets on AVD with *stop* action.

Table 3: Comparing navigation performance (success rate and SPL) on novel scenes from AI2-THOR with *stop* action. Note that all methods in this comparison use RGB input.

Category	Kitchen	Living	Bed.	Bath.	Avg.
Random Walk	7.0 / 3.5	1.8 / 1.0	2.6 / 1.5	17.9 / 8.0	7.3 / 3.5
TD-A3C-U	8.2 / 3.0	2.0 / 1.1	3.2 / 1.9	19.0 / 9.1	8.1 / 3.0
TD-A3C	11.4 / 1.6	5.6 / 0.4	5.3 / 0.7	24.3 / 2.3	11.7 / 1.3
Gated-LSTM-A3C	13.1 / 3.2	4.9 / 1.1	5.1 / 1.2	19.3 / 7.9	10.6 / 3.4
I2A	12.3 / 1.9	5.4 / 0.3	6.2 / 0.8	22.3 / 2.1	11.5 / 1.3
Ours	19.8 / 10.6	11.5 / 5.3	13.6 / 5.9	21.9 / 9.6	16.7 / 7.9

target points can be found in the supplemental material. The plot in Figure 3(left) compares average success rate (without the *stop* action) of different models over an increasing number of time steps, tested on AVD. Our method achieves the steepest increase.

Table 3 evaluates target-driven navigation over synthetic scenes from AI2-THOR. For each of the four room categories, 1000 randomly generated navigation tasks are sampled from the testing split of the dataset. All methods being compared take RGB input, following the original work (Zhu et al. 2017). The random walk baseline can be used as a reference to assess the difficulty of the navigation tasks. For example, living rooms are more challenging while small bathrooms are relatively easy. For the bathrooms, however, our method fails to beat the TD-A3C, because the transparent glass and texture-less furniture make it difficult for our model to infer the surrounding layout which is important to NEO imagination and action prediction. Overall, our model has much better cross-scene generality.

Ablation study. The lower part of Table 2 shows an ablation study. Comparing to the front-view only input, four-view input leads to better results. Our generative method performs much better than its non-generative variant (NoGen) under the same amount of training. This conforms to the consensus that learning a stochastic state space is often more data-efficient than learning a deterministic one (Buesing et al. 2018). The comparison to NoMoP shows that our carefully designed mixture-of-posteriors prior leads to a more powerful internal model by overcoming the over-regularization of latent space caused by the commonly used standard normal distribution prior.

Table 4: Performance (success rate and SPL) for different number of training scenes from AVD without *stop* action.

# Scenes	8	6	4	2
# Samples	616, 630	524, 934	313, 652	152, 836
TD-A3C-U	26.0 / 9.5	25.9 / 8.3	25.9 / 8.5	23.4 / 6.9
TD-A3C	33.4 / 6.1	32.1 / 5.9	28.4 / 3.9	25.9 / 2.8
Gated-LSTM-A3C	25.1 / 9.4	24.7 / 8.9	20.4 / 6.0	19.3 / 4.8
I2A	31.4 / 7.1	29.7 / 6.9	28.2 / 3.2	26.1 / 2.5
Ours	47.9 / 25.8	47.1 / 24.9	45.3 / 22.7	35.1 / 16.9

Table 5: Performance (success rate) comparison of semantic-driven navigation on the AVD test split.

Target label	Couch	Table	Refrig.	Micro.	TV	Avg.
TD-Semantic (Obj.)	80.0	38.0	68.0	38.0	44.0	53.6
Ours (RGB)	64.7	73.7	61.3	38.7	31.3	53.9
Ours (Depth)	83.4	67.4	57.8	41.1	82.0	66.3
Ours (RGBD)	73.5	83.5	72.1	41.5	11.8	56.4

Input modality. Through comparison on different input modalities (Tables 2), the conclusion is that depth information is apparently more useful to our model. This is because depth images contain rich geometry information which benefits a powerful reasoning about the surrounding layout and the modeling of action-observation dynamics.

Data-efficiency. We also evaluate how well our model generalizes when trained on decreasing numbers of scenes (or training samples) from the training split of AVD; see Table 4. The evaluation involves 1000 different navigation tasks sampled from the testing split of AVD. All models show increasing success rates and SPLs with increasing numbers of training scenes. From the results, our method performs consistently better than the alternatives, demonstrating better data-efficiency. In the supplemental material, we compare the training curves of the methods.

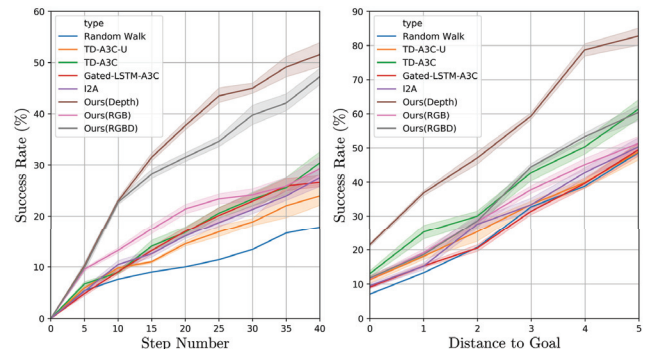


Figure 3: Left: Success rate over increasing number of time steps. Right: Success rate over different values of distance-to-goal thresholds. Each curve is measured based on 1000 navigation tasks from the AVD test split.

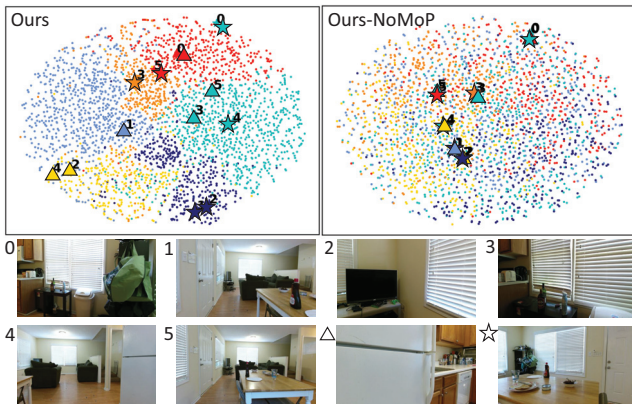


Figure 4: A t-SNE visualization of the latent space $z \sim q(z|x, g)$ of our model (top-left) and its NoMoP variant (top-right). The color of data points indicates action prediction. Some of the data points are marked with an index of the current (front-view) observation x and a shape symbol indicating the target view g (see the indexing of the corresponding view images at the bottom). The ground-truth action of a data point is visualized as the fill color of the corresponding shape symbol. From the color correspondence between the data points and the co-located shape symbols, our model leads to more accurate action prediction.

Close-to-goal stability. In most navigation methods, the agent’s path tends to oscillate when the agent gets close to the goal. The main reason is that situations in which the agent is close to the goal are generally sparse in training. This leads to imbalanced positive and negative situations in training data. Therefore, it is difficult for the trained agent to make a *stop* decision precisely and decisively when approaching the goal. Using 1000 sampled navigation tasks, Figure 3(right) studies the percentage of tasks that succeeds within 40 time steps (disabling the *stop* action) over varying distance-to-goal thresholds used for judging navigation success. In general, smaller thresholds lead to lower success rates due to a higher chance of close-to-goal oscillation. The plot shows that our method achieves more stable close-to-goal convergence for all thresholds, thanks to the expressive approximation of variational inference learned through modeling the agent-environment interaction. The latter leads to high data efficiency even for sparse training samples.

Navigation driven by semantic labels. In methods like TD-Semantic (Mousavian et al. 2018), the navigation goal is defined in the form of a one-hot vector over a prescribed set of semantic labels; for example, $\{Couch, Table, Refrigerator, Microwave, TV\}$. To compare with TD-Semantic, we adapt our method to take the same navigation goal. The comparison is conducted on AVD with the same training/testing split, and the success criterion is within 5 steps to the goal, as in (Mousavian et al. 2018). TD-Semantic can learn visual representations for navigation either from RGB and/or depth input or from semantic input of object detection and segmentation obtained by off-the-

shelf state-of-the-art methods. Under the same input modality, our method outperforms TD-Semantic by 23% for RGB input, 35% for depth input, and 28% for RGBD input for average success rate. Our best performance (with depth input) is 12.7% higher than theirs with semantic input.

Table 5 reports the breakdown results over various target labels for TD-Semantic with semantic input and our method with RGB and/or depth as input. We attribute the good performance to the natural design of the learning task in our model. In TD-Semantic, a deep neural network is learned to predict action cost from the current observation, the goal and the previous action. In contrast, our model predicts the next observation from a latent space modeling the dynamics of action-driven observation transition, making it easier to learn an enriched, meaningful representation (see Figure 4). Moreover, the variational inference module learns to reason about the surrounding layout based on multi-view observations, which is helpful for goal-directed decision making even if the goal is represented in an abstract form of a semantic label instead of a view image.

Visualization of the latent space. To investigate how well our latent space models the navigation policy based on the current observation and the target view, we show in Figure 4 a t-SNE visualization of the latent space $z \sim q(z|x, g)$ learned by our model and its NoMoP baseline. There are two observations. First, the latent space of our model exhibits clear structure w.r.t. action predictions (see the color-coding), making it well suited for navigation decision making. Such expressive latent distribution is facilitated by imposing the mixture-of-posteriors prior conditioned on current observations and next actions $p(z|x, a)$. In contrast, the latent space constrained by a standard Gaussian prior (NoMoP) $p(z)$ is highly unstructured. Second, our action prediction is highly accurate (see the correspondence between data point color (action prediction) and shape symbol fill color (ground-truth action)) thanks to the action-driven variational model for NEO estimation and the separate capacity in the network of action prediction.

In the supplemental material, we also provide visualization of the agent paths for different navigation tasks in unseen scenes, as well as visual comparison of navigation paths against alternative methods.

5 Conclusion

We have presented a generative model for visual navigation that predicts the next action based on the imagination of the next expected observation. The NEO generation models the forward dynamics of agent-environment interaction. The expressive approximation of the variational posterior as a Mixture of Gaussians leads to a data-efficient model with strong model generality. We see great potential in incorporating this generative model into a deep RL framework to address the model imperfection issue in novel scenes.

Acknowledgments

We thank Xingyu Xie for fruitful discussions in the early stage of this project. This work was supported in part

by grants from ARO (W911NF-19-1-0069) and NSFC (61772267, 61572507, 61532003, 61622212).

References

- Ammirato, P.; Poirson, P.; Park, E.; Košecká, J.; and Berg, A. C. 2017. A dataset for developing and benchmarking active vision. In *Proc. ICRA*, 1378–1385.
- Anderson, P.; Chang, A.; Chaplot, D. S.; Dosovitskiy, A.; Gupta, S.; Koltun, V.; Kosecka, J.; Malik, J.; Mottaghi, R.; Savva, M.; et al. 2018. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*.
- Buesing, L.; Weber, T.; Racaniere, S.; Eslami, S.; Rezende, D.; Reichert, D. P.; Viola, F.; Besse, F.; Gregor, K.; Hassabis, D.; et al. 2018. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*.
- Cremer, C.; Li, X.; and Duvenaud, D. 2018. Inference suboptimality in variational autoencoders. In *Proc. ICML*.
- Deisenroth, M., and Rasmussen, C. E. 2011. Pilco: A model-based and data-efficient approach to policy search. In *ICML*, 465–472.
- Dosovitskiy, A., and Koltun, V. 2016. Learning to act by predicting the future. *arXiv preprint arXiv:1611.01779*.
- Finn, C., and Levine, S. 2017. Deep visual foresight for planning robot motion. In *Proc. ICRA*, 2786–2793.
- Gordon, D.; Kadian, A.; Parikh, D.; Hoffman, J.; and Batra, D. 2019. Splitnet: Sim2sim and task2task transfer for embodied visual navigation. *arXiv preprint arXiv:1905.07512*.
- Gupta, S.; Davidson, J.; Levine, S.; Sukthankar, R.; and Malik, J. 2017. Cognitive mapping and planning for visual navigation. In *Proc. CVPR*, 2616–2625.
- Henriques, J. F., and Vedaldi, A. 2018. Mapnet: An allocentric spatial memory for mapping environments. In *CVPR*, 8476–8484.
- Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2016. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lee, L.; Parisotto, E.; Chaplot, D. S.; Xing, E.; and Salakhutdinov, R. 2018. Gated path planning networks. *arXiv preprint arXiv:1806.06408*.
- Leibfried, F.; Kushman, N.; and Hofmann, K. 2016. A deep learning approach for joint video frame and reward prediction in atari games. *arXiv preprint arXiv:1611.07078*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- Marco, A.; Berkenkamp, F.; Hennig, P.; Schoellig, A. P.; Krause, A.; Schaal, S.; and Trimpe, S. 2017. Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization. In *Proc. ICRA*, 1557–1563.
- Mirowski, P.; Pascanu, R.; Viola, F.; Soyer, H.; Ballard, A. J.; Banino, A.; Denil, M.; Goroshin, R.; Sifre, L.; Kavukcuoglu, K.; et al. 2016. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*.
- Mirowski, P.; Grimes, M.; Malinowski, M.; Hermann, K. M.; Anderson, K.; Teplyashin, D.; Simonyan, K.; Zisserman, A.; Hadsell, R.; et al. 2018. Learning to navigate in cities without a map. In *Proc. NeurIPS*, 2419–2430.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *Proc. ICML*, 1928–1937.
- Mousavian, A.; Toshev, A.; Fiser, M.; Kosecka, J.; and Davidson, J. 2018. Visual representations for semantic target driven navigation. *arXiv preprint arXiv:1805.06066*.
- Oh, J.; Guo, X.; Lee, H.; Lewis, R. L.; and Singh, S. 2015. Action-conditional video prediction using deep networks in atari games. In *Proc. NeurIPS*, 2863–2871.
- Oh, J.; Singh, S.; and Lee, H. 2017. Value prediction network. In *Proc. NeurIPS*, 6118–6128.
- Pascanu, R.; Li, Y.; Vinyals, O.; Heess, N.; Buesing, L.; Racaniere, S.; Reichert, D.; Weber, T.; Wierstra, D.; and Battaglia, P. 2017. Learning model-based planning from scratch. *arXiv preprint arXiv:1707.06170*.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 16–17.
- Racaniere, S.; Weber, T.; Reichert, D.; Buesing, L.; Guez, A.; Rezende, D. J.; Badia, A. P.; Vinyals, O.; Heess, N.; Li, Y.; et al. 2017. Imagination-augmented agents for deep reinforcement learning. In *Proc. NeurIPS*, 5690–5701.
- Savinov, N.; Dosovitskiy, A.; and Koltun, V. 2018. Semi-parametric topological memory for navigation. *arXiv preprint arXiv:1803.00653*.
- Savva, M.; Kadian, A.; Maksymets, O.; Zhao, Y.; Wijmans, E.; Jain, B.; Straub, J.; Liu, J.; Koltun, V.; Malik, J.; et al. 2019. Habitat: A platform for embodied ai research. *arXiv preprint arXiv:1904.01201*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International Conference on Machine Learning*, 1889–1897.
- Silver, D.; van Hasselt, H.; Hessel, M.; Schaul, T.; Guez, A.; Harley, T.; Dulac-Arnold, G.; Reichert, D.; Rabinowitz, N.; Barreto, A.; et al. 2017. The predictor: End-to-end learning and planning. In *Proc. ICML*, 3191–3199.
- Tamar, A.; Wu, Y.; Thomas, G.; Levine, S.; and Abbeel, P. 2016. Value iteration networks. In *Proc. NeurIPS*, 2154–2162.
- Tomczak, J. M., and Welling, M. 2017. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*.
- Watter, M.; Springenberg, J.; Boedecker, J.; and Riedmiller, M. 2015. Embed to control: A locally linear latent dynamics model for control from raw images. In *Proc. NeurIPS*, 2746–2754.
- Wu, Y.; Wu, Y.; Gkioxari, G.; and Tian, Y. 2018. Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209*.
- Yang, W.; Wang, X.; Farhadi, A.; Gupta, A.; and Mottaghi, R. 2018. Visual semantic navigation using scene priors. *arXiv preprint arXiv:1810.06543*.
- Zhang, J.; Springenberg, J. T.; Boedecker, J.; and Burgard, W. 2017. Deep reinforcement learning with successor features for navigation across similar environments. In *Proc. IROS*, 2371–2378.
- Zhu, Y.; Mottaghi, R.; Kolve, E.; Lim, J. J.; Gupta, A.; Fei-Fei, L.; and Farhadi, A. 2017. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *ICRA*, 3357–3364.