

Symmetry-aware Neural Architecture for Embodied Visual Exploration

Shuang Liu
RIKEN Center for AIP
shuang.liu.ej@riken.jp

Takayuki Okatani
Tohoku University, RIKEN Center for AIP
okatani@vision.is.tohoku.ac.jp

Abstract

Visual exploration is a task that seeks to visit all the navigable areas of an environment as quickly as possible. The existing methods employ deep reinforcement learning (RL) as the standard tool for the task. However, they tend to be vulnerable to statistical shifts between the training and test data, resulting in poor generalization over novel environments that are out-of-distribution (OOD) from the training data. In this paper, we attempt to improve the generalization ability by utilizing the inductive biases available for the task. Employing the active neural SLAM (ANS) that learns exploration policies with the advantage actor-critic (A2C) method as the base framework, we first point out that the mappings represented by the actor and the critic should satisfy specific symmetries. We then propose a network design for the actor and the critic to inherently attain these symmetries. Specifically, we use G-convolution instead of the standard convolution and insert the semi-global polar pooling (SGPP) layer, which we newly design in this study, in the last section of the critic network. Experimental results show that our method increases area coverage by $8.1m^2$ when trained on the Gibson dataset and tested on the Matterport3D dataset, establishing the new state-of-the-art.

1. Introduction

Embodied visual navigation is one of the key problems of autonomous navigation and has attracted increasing attention recently. Researchers have studied various target-oriented tasks of visual navigation so far. These include point goal navigation aiming to reach a given coordinate in an environment [51], object goal navigation aiming to find a specific object in an environment [7,20,54], image goal navigation aiming to navigate to a location in an environment designated by an image [36,39], reconstruction navigation deciding where to look next to reconstruct a scene [21,38], and room goal navigation aiming to navigate to a specified room [47].

In contrast with these tasks, visual exploration is universal and task-agnostic; an agent is asked to visit all the nav-

igable areas of an unknown environment as quickly as possible. Thus, it can serve as an upstream task for the above target-oriented tasks. For instance, the agent may first explore the unseen environment and establish the knowledge about the environment, after which the agent utilizes it to perform specific tasks efficiently [36].

Numerous methods have been proposed for the problem of visual exploration so far. These methods are divided into two categories, heuristic methods [14,44,49] and learning-based methods. Recent studies have mainly considered learning-based methods.

Visual exploration is a problem of choosing the optimal action at each time step to maximize the above exploration goal, given a history of inputs to the agent. Recent studies employ simulators of virtual environments, such as Habitat [42], to consider the problem under realistic conditions. In such simulators, photorealistic visual inputs from virtual environments are available. To predict the optimal action from the agents' input, including the visual inputs, learning-based methods typically use reinforcement learning (RL) to acquire the mapping between the two from an enormous amount of agents' interaction with the virtual environments.

Although they have achieved reasonably good exploration performance, the learning-based methods are inherently vulnerable to the statistical deviations between training and test data. Specifically, a model trained on a set of training data works well on novel environments that are statistically similar to the training data but yields suboptimal performance on novel environments that are dissimilar, i.e., out-of-distribution (OOD) inputs. This is not the case with the heuristic methods.

To apply the learned models to real-world environments, it is necessary to increase their generalization ability in the above sense. How is it possible? A promising approach is data augmentation, and it has been widely adopted in both deep learning [13,24,52] and reinforcement learning [22,33,50]. However, this approach has its limitations.

In this study, we pay attention to the inductive biases specific to the task. Using them, we want to narrow the solution space for RL and thus avoid overfitting to training data, aiming to improve the above generalization ability. Specif-

ically, we employ the framework of active neural SLAM (ANS) [6], which is the most successful for the task, and consider how to do the above within the framework.

ANS is a method with a modular and hierarchical structure. First, a module yields a long-term goal from a history of the agent’s inputs. The mapping between them is called global policy, which is represented by a neural network. Next, another module computes the shortest path from the agent’s position to the given long-term goal. The path is then subsampled to construct a sequence of short-term goals. Selecting a short-term goal in turn, another module generates actions to reach the short-term goal.

The module that plays a key role in ANS is the neural network that represents the global policy. (The problems solved by the other modules are not so hard.) ANS employs the advantage actor-critic (A2C) method [28] to learn the global policy. The actor and the critic are modeled by neural networks, specifically, a pair of independent networks or a single network having two output branches. The actor and the critic receive as their input the current estimate of the environment’s map along with the agent’s navigation paths etc. expressed in the form of a multi-channel 2D map. The actor yields a 2D likelihood map for the long-term goal, whereas the critic predicts the future accumulated reward as a scalar value.

In this paper, we first point out that specific symmetries should exist for each mapping represented by the actor and the critic. Concretely, the 2D map predicted by the actor should be equivariant to translation and rotation of its input having the form of a 2D map. The reward predicted by the critic should be invariant to input rotation and should not be invariant to input translation. However, standard convolutional neural networks (CNNs), employed in ANS, do not by nature have these symmetries.

To cope with this, we design the network to inherently possess the desired symmetries for the actor and the critic. Specifically, we redesign the network of ANS as follows. First, we employ $p4$ G -convolution [12], which approximately achieves rotation equivariance, for all the convolutional layers and replace the max pooling with blur pooling [53], which achieves translation equivariance more accurately. Then, the resulting network representing the actor becomes equivariant to input translation and approximately to input rotation. For the critic, employing the same stack of convolutional layers (rigorously, it is shared by the actor), we design and place a new building block named a semi-global polar pooling (SGPP) layer on top of the convolutional block. The SGPP layer applies polar mapping to its input and then pooling in the circumferential direction in the input map. Then, the resulting critic becomes not invariant to input translation and approximately invariant to input rotation.

We conduct experiments to test our approach using the

Gibson [48] and the Matterport3D (MP3D) [5] datasets. The results show that the proposed method improves the original ANS by a good margin, proving the approach’s effectiveness. Specifically, it improves ANS by $8.1 m^2$ (from 76.3 to $84.4 m^2$) in area coverage, the primary metric for evaluating exploration performance, in the setting to evaluate the generalization ability, i.e., training the models on Gibson and testing them on MP3D.

2. Related Works

2.1. Visual Navigation

Classical navigation methods leverage simultaneous localization and mapping (SLAM) to construct environmental maps, including occupancy map [4] and topological map [36]. Agents navigate to points in the environments. This is known as map-based navigation [2].

In contrast with the map-based navigation that assumes a known map, target-driven navigation tasks have recently emerged, which are termed mapless navigation. In these tasks, the movement of the agent is determined by visual cues provided by the environment. Example tasks are object-goal, room-goal, and image-goal navigation.

Object-goal navigation requires agents to navigate to designated objects in an unknown environment. Either implicitly learned [54] or explicitly encoded [17, 25] relationship of objects is utilized to facilitate finding the objects. Room-goal navigation requires agents to arrive at a specific room as quickly as possible. A method is proposed in [47] to learn a probabilistic relation graph to acquire prior knowledge about the layout of environments.

Image-goal navigation asks agents to navigate to a location in an unknown environment that is specified by its image. A popular solution is to take both current and goal observation as inputs and employ a Siamese network to perform the navigation effectively. The method proposed in [26] combines an attention mechanism with the Siamese network for building the memory of environments, which is then used by the policy for image-driven navigation. The work [11] follows the same formulation based on a Siamese network but exploits information obtained through keypoint matching, generating a self-supervised reward.

2.2. Visual Exploration

Visual exploration has received considerable attention due to its task-agnostic nature [6, 9, 21, 27, 34]. Classical approaches continually select vantage points, such as frontier points [49], to visit. The downside of these approaches is that they are vulnerable to external noises such as sensor noises.

The recent boom of deep reinforcement learning (RL) has created a new wave of development of visual exploration. Numerous works have casted the visual exploration

as a partially observable Markov decision process and solve it under the framework of RL [1, 6, 9, 21, 27, 30–32, 34, 35]. Most of them use an actor-critic architecture [1, 6, 30, 31, 34, 35].

The actor learns a navigation policy, and the critic estimates a value function given the agent’s current state. Existing approaches to improve performance have been geared towards novel designs of reward, such as coverage, curiosity, and novelty [35], or novel architecture design, e.g., the hierarchical architecture [6]. There is a study [9] that proposes to utilize area coverage reward along with supervision from human experience to learn policies.

An excellent standard framework, which is actor-critic based as well, was constructed by [35] to evaluate multiple rewards for exploration. The approaches above all solve the exploration problem in an end-to-end manner, directly mapping visual data to actions.

In contrast to this, the work [6] tackled it hierarchically. It divided the mapping into two steps. First, a global target point, which was similar to vantage points in classical exploration methods, was computed by a global policy network. Then a local policy network produces actions to reach the waypoints generated by the path planner according to the global target point. Occupancy anticipation [34] also predicts and rewards invisible spatial maps to improve map accuracy.

So far, existing methods have produced impressive results in the case where the testing environments are similar to training environments in terms of layout, area, etc. Nevertheless, they still struggle with the cases where training and testing environments are different, hindering their deployment in practice. It is primarily because RL notoriously struggles with (poor data efficiency and) generalization capabilities. We aim to enhance generalization ability for visual navigation by injecting inductive bias about symmetry into networks.

2.3. Equivariance and Invariance

Recent years have witnessed the great success of convolutional neural networks (CNNs) in computer vision tasks. CNNs have built-in translation-equivariance, and those with downsampling/pooling operations have built-in local translation invariance. These contribute to the successes in applications to various problems. It could be possible to learn these symmetries from data if the model has sufficient parameters. However, it will increase the risk of overfitting.

A function is equivariant if the output changes in the same way that the input changes [19]. Novel convolution layers with different equivariance have been proposed so far, group equivariant convolution networks [12, 16], steerable convolution networks and harmonic networks for rotation equivariance, scale equivariance [23, 46], and permutation equivariance [43]. Such convolution layers equipped

with various types of equivariance have been proven to be beneficial to better performance in tracking [40], classification, trajectory prediction [45], segmentation [29], and image generation [15]. However, these methods have never been applied to visual navigation.

A primary purpose of pursuing various types of equivariance is to achieve a certain type of invariance. A mapping is called invariant if the output remains the same no matter how the input changes. Although data augmentation and invariance-oriented loss functions [10] may enhance global invariance, it is not guaranteed to generalize to out-of-distribution data. By contrast, global invariance can be imposed by a global pooling layer following equivariant layers. Global rotation invariance is enforced in [10] for texture classification by combining equivariant convolution layers with a global average pooling layer in this order. Several studies [23, 41] attain global scale-invariance by combining layers having scale-equivariance with global max pooling.

3. Symmetry-aware Neural Architecture

3.1. Problem of Visual Exploration

Following previous studies [6, 9, 34], we consider a visual exploration task in which an agent explores an unknown 3D environment, e.g., a floor of a building, that allows only 2D motion. The goal for the agent is to go everywhere in the environment it can go while creating an environment’s 2D map. We employ the Habitat simulator [37] as the framework for studying the task.

The agent receives several inputs from the environment, i.e., an image or a depth map of the scene in front of it, the agent’s pose, and an odometry signal measuring its motion. Then, the agent moves to explore the environment by feeding the actuation signal to its motor. Thus, the problem is to compute the actuation signal and update the environment map at every time step, given the history of the inputs the agent has received until then. See the literature [6] for more details.

Recent studies have formulated the problem as learning policies that yield actions (or intermediate representation leading to actions) from the received inputs; most of them employ reinforcement learning (RL) [6, 9, 21, 34]. Recently, Chaprot et al. [6] proposed a method named active neural SLAM (ANS), having established the new state-of-the-art. Our study is built upon ANS and we will first revisit and summarize it below.

3.2. Revisiting Active Neural SLAM (ANS)

ANS has a modular and hierarchical structure to better solve the visual exploration problem. Instead of learning a direct mapping from the inputs to an action, ANS learns two different policies in a hierarchy, i.e., a global policy that

yields a long-term goal in the 2D environment map and a local policy that yields actions to approach a short-term goal, which is subsampled from the path to the long-term goal. More specifically, ANS consists of four modules, i.e., a neural SLAM module, a global policy module, a path planner, and a local policy module.

The neural SLAM module computes a local top-view 2D egocentric map $p_t \in [0, 1]^{2 \times v \times v}$ and estimates an accurate agent pose from the current inputs. (The first and second channel of p_t represents the obstacle and the explored region at time t , respectively.) Then, the computed local egocentric map is registered to the global map $h_t \in [0, 1]^{4 \times M \times M}$ using the estimated pose. (The first and second channel of h_t represents the obstacles and explored area at time t respectively. The third channel indicates the computed agent position. The last channel records the path that has been visited by the agent.)

The global policy module makes ANS the most distinct from other methods. It predicts a long-term goal g_t given two different views h_t^g and h_t^l of the latest map h_t ; h_t^g is a rescaled version of h_t , and h_t^l is a local view cropped from h_t with the agent position as its center; and h_t^g and h_t^l have the same size $G \times G$ and are concatenated in the channel dimension.

The module is implemented as an advantage actor-critic (A2C) network [28]. As shown in Fig. 1, it consists of an actor network and a critic network. Each network has a similar design of a convolutional block, which consists of a stack of convolutional and max-pooling layers, and subsequent fully connected layers. In a standard design, they share the convolutional block.

The actor network outputs a 2D likelihood map of the long-term goal g_t represented in the coordinates of the global map h_t^g . It learns a global policy, denoted by $g_t = \pi(s_t | \theta_G)$, where $s_t = (h_t^l, h_t^g)$ is the state at time t and θ_G are parameters. The critic network represents a value function $V(s_t | \theta_V)$; $V(s_t | \theta_V)$ estimates the expected future accumulated reward with the agent, which is currently at state s_t and will take actions by following the current policy $\pi(s_t | \theta_G)$. The parameters θ_G and θ_V are optimized to maximize the area coverage.

The other two modules, i.e., the path planner and the local policy module, play the following roles. The path planner computes the shortest path from the current location to the long-term goal, which is subsampled to generate a number of short-term goals. The local policy module predicts actions to reach the next short-term goal.

3.3. Symmetries in the Global Policy

3.3.1 Outline

Although ANS performs reasonably well and is currently state-of-the-art, its network is a plain CNN consisting only of generic network components, such as convolutional and

fully connected layers. Thus, there is room for improvement in architectural design, and we aim to find a design that better fits the problem structure.

Toward this end, sticking to the basic framework of ANS that is proven to be effective, we propose to redesign its global policy module, which consists of the actor and critic networks. We consider what conditions the mappings realized by the two networks should satisfy. We pay attention to symmetry of the mappings, more specifically, their equivariance and invariance to translation, rotation, and scaling. Leveraging the fact that the exploration task is a geometric problem, we derive what symmetry the ideal actor and critic networks should implement, based on which we redesign the two networks.

In other words, we incorporate the inductive bias of the task into the network design. A neural network represents a mapping, and we train it using data to hopefully find an optimal mapping. The network’s architecture specifies a subspace in the mapping space. Enforcing the network to inherently satisfy a required constraint by its design shrinks the subspace, leading to a better solution.

3.3.2 Invariance and Equivariance

As mentioned above, we are interested in *equivariance* and *invariance*. Their mathematical definitions are as follows. Given a group \mathcal{G} on homogeneous spaces \mathcal{X}_1 and \mathcal{X}_2 , a mapping $\Psi : f(\mathcal{X}_1) \rightarrow f'(\mathcal{X}_2)$ that has the property

$$\Psi[\mathbb{T}_g f(x)] = \mathbb{T}'_g \Psi[f(x)], \forall g \in \mathcal{G}, \quad (1)$$

is said to be *equivariant* to the group \mathcal{G} (or \mathcal{G} -equivariant) and to be *invariant* to the group \mathcal{G} (or \mathcal{G} -invariant) if $\mathbb{T}'_g = e$, the identity [3]. \mathbb{T}_g is the transformation corresponding to its group action g .

Intuitively, they are interpreted as follows. Suppose a mapping that receives an input and yields an output and also applying a geometric transformation to the input of the mapping. If the mapping is invariant to the transformation, the output for the transformed input will remain the same as the original output. If the mapping is equivariant to the transformation, that will be the same as the result of applying the same transformation to the original output.

Convolutional layer are equivariant to input translation. Thus, shifting an input map results in the same shift of the output map. This applies to the input/output of a stack of any number of convolutional layers. (Rigorously, the output map undergoes the input shift plus downscaling corresponding to the downsampling in those layers.) On the other hand, global average pooling (GAP) layers are invariant to input translation. Thus, CNNs having a GAP layer on top of the stack of convolution layers are invariant to input translation. CNNs and a stack of convolution layers are not equivariant to rotation nor scaling. They are not invariant to them, either.

3.3.3 Ideal Symmetries of Actor and Critic Networks

Now we consider what symmetry the actor and critic networks should have. First, the actor network receives a 2D map encoding the current state of the agent, and outputs a 2D map containing a long-term goal, as explained above and shown in Fig. 1. It is easy to see that this mapping should ideally be equivariant to translation and rotation. It should not be equivariant to scaling.

The critic network receives the same input map and outputs a scalar value, the future accumulated reward when following the current policy given the current state. The mapping should be invariant to rotation since the future accumulated reward should be independent of the orientation of the input map. It should not be invariant to translation or scaling, since applying these transformation to the input map should change the reward.

Table 1 (the first two rows) summarizes the requirements for the ideal actor and critic networks.

To design a new critic network, we limit our attention to networks having a stack of convolutional layers at their initial section. This is also the case with the critic network of ANS, which has a stack of fully connected layers right after the convolutional layer stack, and we will redesign this section. Now, for the mapping from the input to the final output to be invariant to rotation as mentioned above, the mapping to the last convolutional layer needs to be equivariant to rotation. This is because otherwise, it will be extremely hard, if not impossible, to achieve the rotation-invariance at the final output. We express this requirement in the ‘↑ at the last conv.’ row in Table 1.

The actor and critic networks of the original ANS does not have the same symmetries as the ideal ones. The missing symmetries are indicated by the \times ’s in Table 1. Specifically, the actor should be rotation-equivariant but is not; the critic should be rotation-invariant but is not. As discussed above, the mapping from the input to the last convolutional layer of the critic should be rotation-equivariant but is not.

3.4. Rotation-Equivariance of the Actor Network

Thus, we need to newly equip the actor network with rotation-equivariance while maintaining its translation-equivariance. To do so, we propose to employ G -convolution [12] and blur pooling [53]; specifically, we replace the standard convolution and max pooling of the original ANS actor network with the two, respectively.

Considering computational efficiency, we employ $p4$ G -convolution. All combinations of translations and rotations by 90 degrees form the group $p4$, which can be parameterized by

$$g(m, z_1, z_2) = \begin{bmatrix} \cos(m\pi/2) & -\sin(m\pi/2) & z_1 \\ \sin(m\pi/2) & \cos(m\pi/2) & z_2 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2)$$

	Equivariance			Invariance		
	Trans.	Rot.	Scale	Trans.	Rot.	Scale
Ideal actor	✓	✓	✗			
Ideal critic				✗	✓	✗
↑ at the last conv.	–	✓	–			
ANS actor	✓	✗	✗			
ANS critic				✗	✗	✗
↑ at the last conv.	✓	✗	✗			

Table 1. Upper three rows: ideal symmetries that should be implemented by the actor and critic networks of the ANS framework. Lower three rows: actual symmetries realized by the original actor and critic networks of ANS. ✓ indicates the network is equipped with the symmetry. ✗ indicates the network does not have the symmetry. ‘–’ indicates the symmetry is not specified. Blank cells mean the symmetry is irrelevant. ✗ in red color indicates that the symmetry of ANS networks differs from the ideal one.

where $m \in \{0, 1, 2, 3\}$ and $(z_1, z_2) \in \mathbb{Z}^2$ [12].

As shown in Fig. 1, $p4$ G -convolution can be performed by first rotating filters with angles $m\pi/2, m = \{0, 1, 2, 3\}$ to form a filter bank, and then applying it to the input feature map.

As mentioned above, we redesign the convolutional layer stack of the actor network with the same number of layers performing $p4$ G -convolution. The original ANS actor network employs max pooling for downsampling the feature map. However, max pooling makes the translation-equivariance inaccurate. Max pooling ignores the Nyquist sampling theorem, breaking translation equivariance. Blur pooling filters the signal before downsampling to better preserve translation equivariance. Hence, we employ blur pooling and replace max pooling with this.

By revising the convolutional block (i.e., the stack of convolutional and pooling layers) in the original ANS actor network as above, the mapping represented by the block becomes approximately equivariant to rotation and precisely equivariant to translation. We need to maintain the fully connected layers after the convolutional block since predicting the agent’s goal seems to need them to integrate the global features in a non-simple manner. Fully connected layers are not inherently equipped with the desired equivariance, and we leave it to training; owing to their flexibility, we expect the actor network to gain the desired equivariance at the final output.

On the other hand, we want to make the critic network have rotation-invariance. As mentioned earlier, to do so, we want the mapping represented from the input to the output of the convolutional block to be rotation-equivariant. We employ here the design of the original ANS that the critic network share the convolutional block with the actor network. Then, the above redesign of the convolutional block achieves what we want. We further revise the section on top of the convolutional block to realize the desired symmetries

for the critic network.

3.5. Rotation-Invariance of the Critic Network: Semi-global Polar Pooling

The critic network should represent a mapping that is rotation-invariant and is not translation-invariant or scale-invariant, as shown in Table 1. We assume the convolutional block to have rotation-equivariance due to its design, as explained above. To attain the above (in)variance, we then propose a new network component, named the semi-global polar pooling (SGPP) layer.

SGPP is illustrated in Fig. 2. It first converts the input feature map from Cartesian space to polar space. Let $I(x, y) \in \mathbb{R}^{c \times h \times w}$ denote the feature map represented in Cartesian space, where x, y are the Cartesian coordinates. The feature map $I'(\rho, \phi) \in \mathbb{R}^{c \times h \times w}$ represented in polar space is given by

$$\begin{aligned} \rho &= \sqrt{x^2 + y^2}, \\ \phi &= \arctan \frac{y}{x}, \end{aligned} \quad (3)$$

where ρ and ϕ are the coordinates in polar space. We finally apply average-pooling to the features over all ϕ 's as

$$\Phi = \frac{1}{n} \sum_{\phi} I'(\rho, \phi), \quad (4)$$

to obtain a pooled feature $\Phi \in \mathbb{R}^{c \times h}$.

As the convolutional block, whose output is the input to the SGPP layer, is equivariant to rotation, Φ is invariant to rotation since it is pooled over the coordinate ϕ (i.e., the circumferential direction). Φ is further processed by fully connected layers, computing the final output of the critic network. As Φ is already invariant to rotation, these additional layers do not change the invariance. While the convolutional block is also equivariant to translation, the Cartesian-polar conversion invalidates the equivariance; thus Φ is not invariant to translation. It is not invariant to scaling, either. Thus, the critic network having the SGPP layer in between the convolutional block and the fully connected layers attains the desired (in)variance to translation, rotation, and scaling, as illustrated in Table 1.

4. Experimental Results

4.1. Experimental Setup

As with previous studies [6–9, 18, 34, 39], we use the Habitat simulator [37] for our experiments. We choose the configuration that depth images are available as the visual input and the actuation and sensory signals include noises. This applies to the training and test times. To evaluate the performance of exploration methods, we use the area coverage (i.e., the area seen during exploration) within a time

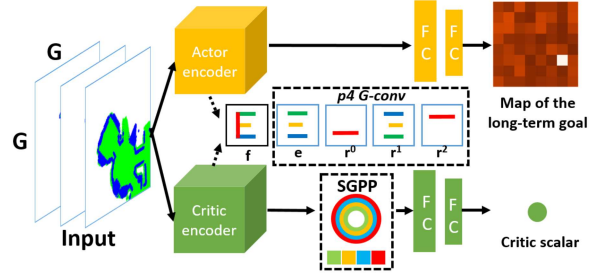


Figure 1. Overall architecture of Symmetric Global Policy Network (using G -convolution, blur pooling and SGPP) and Global Policy Network of ANS (using convolution, max pooling and without using SGPP). The dashed rectangle illustrates the feature maps produced by $p4$ G -convolution assuming its input is \mathbf{E} and the filter is a horizontal edge extractor.




Figure 2. Illustration of the proposed semi-global polar pooling (SGPP).

budget for the primary metric, following [6, 9]. We compare our method with several baselines. We run each method five times and report their average area coverage with the standard deviation. Our code is publicly available ¹.

4.2. Datasets

Following previous studies, we employ two datasets, Gibson [48] and MP3D [5]. Both of them contain photorealistic virtual environments created from real-world scenes. Most environments contained in Gibson are office spaces, while those in MP3D are homes. Thus, the environments of Gibson differ from those of MP3D in terms of scenarios, layout, area, and so on. Overall, MP3D environments are larger in the area and more irregular in the layout. Fig. 3 provides the layouts of several representative examples from Gibson and MP3D. It is seen that there is a gap in various aspects between the two.

To evaluate the generalization ability of the methods across different datasets, we train each model on Gibson and test it either on Gibson (in-distribution) or on MP3D (out-of-distribution), following the experiments of [6].

¹<https://github.com/vincent341/S-ANS>




Figure 3. Environment layout of samples from (a) Gibson and (b) MP3D dataset.

4.3. Compared Methods

We experimentally compare the proposed method with several baselines for global policy. The first is FBE (Frontier based Exploration) [49]. FBE is a heuristic method that iteratively selects a point at the frontiers (boundaries between explored free region and unexplored region) using various strategies. We follow the strategy implemented in [35], selecting a random point at the longest boundary. We replace the global policy module in ANS with FBE and evaluate its performance.

The second is a variant of FBE, named FBE-RL, which replaces the global policy module in ANS with a RL-based FBE. To be specific, FBE-RL selects a target point at the frontiers using the policy produced by an actor-critic network, instead of a heuristic policy. More details of FBE-RL will be given in the supplementary material.

The third is ANS, for which we use the code² and the settings given in [34]. Our method for generating global policy is integrated to ANS by replacing the original global policy module with ours. We call this model symmetry-aware ANS or S-ANS in what follows. We run each model five times and report their average and standard deviation of the area coverage at time step 1,000.

4.4. Results

Table 2 shows the results. Comparing the methods in the case of training and testing on Gibson, the proposed S-ANS outperforms ANS with the margin of $0.8m^2$ (33.7 vs. $32.9m^2$). The margin becomes more significant when testing the same models on MP3D, i.e., $8.1m^2$ (84.4 vs. $76.3m^2$). This demonstrates the better generalization ability of S-ANS.

Fig. 4 shows the representative examples of their exploration paths of ANS and S-ANS on the same four environments of MP3D. The first and second columns of Fig. 4

	Gibson (m^2)	MP3D (m^2)
FBE	26.5 ± 0.5	69.6 ± 1.9
FBE-RL	28.2 ± 0.3	63.0 ± 3.2
ANS	32.9 ± 0.2	76.3 ± 2.8
S-ANS	33.7 ± 0.2	84.4 ± 1.7

Table 2. Exploration performance (in area coverage, m^2) of different models on Gibson and MP3D. All the models are trained on Gibson.




Figure 4. Exploration paths of ANS (top row) and S-ANS (bottom row) for four representative environments from MP3D. The number at the bottom left of each box indicates the area coverage (m^2). The dissimilarity of the environment from the training data tends to widen the gap between the two methods.

show typical examples for which S-ANS explores a much larger area than ANS; S-ANS explores twice as large area as ANS. As with these examples, S-ANS tends to show better exploration performance for environments that are more dissimilar from those of Gibson. When the environments are similar, their performances tend to be close, as shown in the third column of Fig. 4. These observations further verify the improved generalization ability of S-ANS.

It is seen from Table 2 that FBE and FBE-RL show inferior performance. It is noteworthy that FBE and FBE-RL behave differently for different combinations of the train and test data. While FBE-RL performs better on Gibson by $1.7m^2$ (28.2 vs $26.5m^2$), it performs worse on MP3D by $-6.6m^2$ (63.0 vs $69.6m^2$). We can say that learning better policies improves exploration performance when there is only a little gap between the train and test data. However, it is a double-edged sword; it causes overfitting and leads to inferior exploration ability when there is a gap between the train and test data. This demonstrates the difficulty with learning-based approach to the exploration task.

²github.com/facebookresearch/OccupancyAnticipation

4.5. Ablation Study

As explained earlier, our method imposes translation- and rotation-equivariance on the actor network and (approximate) rotation-invariance on the critic network. The former is implemented by the G -convolution and the blur pooling. The latter is implemented by SGPP in addition to G -convolution and the blur pooling. To examine the effectiveness of each component, we create two variants of the proposed network, named E-ANS and G-ANS.

E-ANS is a model created by removing the SGPP layer from the critic network of S-ANS. Without SGPP, it is not (even approximately) invariant to rotation, translation, or scaling. Its actor network maintains the ideal property of equivariance. G-ANS is an intermediate model between E-ANS and S-ANS; it is created by replacing the SGPP layer with a global average pooling (GAP) layer in the critic network of S-ANS. The added GAP layer makes the critic network invariant to both rotation and translation; recall that the convolutional block before the GAP layer is equivariant to rotation and translation due to the employment of the G -convolution and the blur-pooling. The added translation invariance is superfluous compared with the ideal invariance of the critic network shown in Table 1. In short, considering the excesses and deficiencies of the implemented invariance/equivariance, the expected performance will be S-ANS > E-ANS, G-ANS > ANS.

We train E-ANS and G-ANS on Gibson and test them on Gibson and MP3D in the same way as ANS and S-ANS. Figs. 5(a) and (b) show the explored areas by the four models on the test splits of Gibson and MP3D, respectively. Each solid curve and shadowed area indicates the mean and standard deviation over five runs, respectively, of the area coverage at a different time step.

It is seen from the results that the four models are ranked in the performance as S-ANS > E-ANS > G-ANS > ANS. This matches well with our expectation mentioned above. The results further tell us that the excessive translation-invariance of the critic network in G-ANS does more harm than deficient rot-invariance of that of E-ANS. Overall, these results validate the effectiveness of our approach.

5. Conclusion

Existing learning-based methods for visual exploration struggle with generalization to OOD environments, that is, statistically different environments from those used for training. We believe this is partly because the inductive biases available for the task are not used effectively. Employing the framework of active neural SLAM (ANS), we have shown that the actor and the critic should satisfy specific symmetries with their mappings. We then propose a design of neural networks that inherently possesses the ideal symmetries. Specifically, we propose to use G -convolution

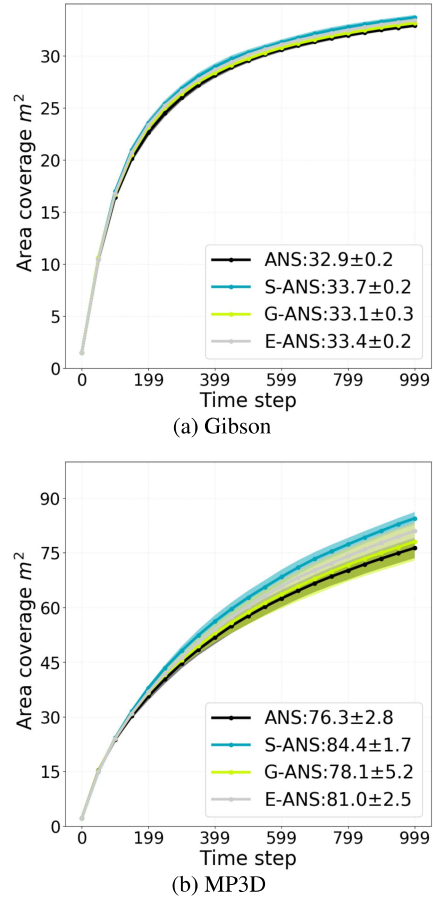


Figure 5. Expolaration performance (in area coverage, m^2) of ANS, S-ANS, G-ANS, and E-ANS tested on (a) Gibson and (b) MP3D. All the models are trained on the Gibson dataset.

instead of the standard convolution layer. We also propose the semi-global polar pooling (SGPP) layer, a network component, that can make the network invariant to rotation and not invariant to translation when using it at the final section of the critic network. Our approach improves the performance of ANS by $8.1m^2$ in area coverage in the setting where training and test datasets are different.

Acknowledgement

This work was partly supported by JSPS KAKENHI Grant Number 20H05952 and 19H01110.

References

- [1] Edward Beeching, Jilles Dibangoye, Olivier Simonin, and Christian Wolf. Egomap: Projective mapping and structured egocentric memory for deep rl. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 525–540. Springer, 2020. 3

- [2] Francisco Bonin-Font, Alberto Ortiz, and Gabriel Oliver. Visual navigation for mobile robots: A survey. *Journal of intelligent and robotic systems*, 53(3):263–296, 2008. [2](#)
- [3] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *CoRR*, abs/2104.13478, 2021. [4](#)
- [4] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on robotics*, 32(6):1309–1332, 2016. [2](#)
- [5] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. [2](#), [6](#)
- [6] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *International Conference on Learning Representations*, 2020. [2](#), [3](#), [6](#)
- [7] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. *Advances in Neural Information Processing Systems*, 33, 2020. [1](#), [6](#)
- [8] Devendra Singh Chaplot, Helen Jiang, Saurabh Gupta, and Abhinav Gupta. Semantic curiosity for active visual learning. In *European Conference on Computer Vision*, pages 309–326. Springer, 2020. [6](#)
- [9] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. In *International Conference on Learning Representations*, 2019. [2](#), [3](#), [6](#)
- [10] Gong Cheng, Peicheng Zhou, and Junwei Han. Learning rotation-invariant convolutional neural networks for object detection in vhr optical remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 54(12):7405–7415, 2016. [3](#)
- [11] Yunho Choi and Songhwa Oh. Image-goal navigation via keypoint-based reinforcement learning. In *2021 18th International Conference on Ubiquitous Robots (UR)*, pages 18–21. IEEE, 2021. [2](#)
- [12] Taco Cohen and Max Welling. Group equivariant convolutional networks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999. PMLR, 20–22 Jun 2016. [2](#), [3](#), [5](#)
- [13] Ekin D. Cubuk, Barret Zoph, Dandelion Mané, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 113–123, 2019. [1](#)
- [14] Anna Dai, Sotiris Papatheodorou, Nils Funk, Dimos Tzoumanikas, and Stefan Leutenegger. Fast frontier-based information-driven autonomous exploration with an mav. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9570–9576, 2020. [1](#)
- [15] Neel Dey, Antong Chen, and Soheil Ghafurian. Group equivariant generative adversarial networks. *CoRR*, abs/2005.01683, 2020. [3](#)
- [16] Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu. Exploiting Cyclic Symmetry in Convolutional Neural Networks. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1889–1898. JMLR, 2016. [3](#)
- [17] Heming Du, Xin Yu, and Liang Zheng. Learning object relation graph and tentative policy for visual navigation. In *European Conference on Computer Vision*, pages 19–34. Springer, 2020. [2](#)
- [18] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9701–9707. IEEE, 2020. [6](#)
- [19] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. [3](#)
- [20] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017. [1](#)
- [21] Dinesh Jayaraman and Kristen Grauman. Learning to look around: Intelligently exploring unseen environments for unknown tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1238–1247, 2018. [1](#), [2](#), [3](#)
- [22] Misha Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *Advances in Neural Information Processing Systems*, 33:19884–19895, 2020. [1](#)
- [23] Tony Lindeberg. Scale-covariant and scale-invariant gaussian derivative networks. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 3–14. Springer, 2021. [3](#)
- [24] Shuang Liu, Mete Ozay, Hongli Xu, Yang Lin, and Takayuki Okatani. A generative model of underwater images for active landmark detection and docking. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8034–8039, 2019. [1](#)
- [25] Yunlian Lv, Ning Xie, Yimin Shi, Zijiao Wang, and Heng Tao Shen. Improving target-driven visual navigation with attention on 3d spatial relationships. *CoRR*, abs/2005.02153, 2020. [2](#)
- [26] Lina Mezghani, Sainbayar Sukhbaatar, Thibaut Lavril, Oleksandr Maksymets, Dhruv Batra, Piotr Bojanowski, and Karteek Alahari. Memory-augmented reinforcement learning for image-goal navigation. *CoRR*, abs/2101.05181, 2021. [2](#)
- [27] Lina Mezghani, Sainbayar Sukhbaatar, Arthur Szlam, Armand Joulin, and Piotr Bojanowski. Learning to visually navigate in photorealistic environments without any supervision. *CoRR*, abs/2004.04954, 2020. [2](#), [3](#)
- [28] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016. [2](#), [4](#)

- [29] Philip Müller, Vladimir Golkov, Valentina Tomassini, and Daniel Cremers. Rotation-equivariant deep learning for diffusion MRI. *CoRR*, abs/2102.06942, 2021. [3](#)
- [30] Tushar Nagarajan and Kristen Grauman. Learning affordance landscapes for interaction exploration in 3d environments. In *Advances in Neural Information Processing Systems*, 2020. [3](#)
- [31] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, pages 2778–2787. PMLR, 2017. [3](#)
- [32] William Qi, Ravi Teja Mullapudi, Saurabh Gupta, and Deva Ramanan. Learning to move with affordance maps. In *International Conference on Learning Representations*, 2020. [3](#)
- [33] Roberta Raileanu, Max Goldstein, Denis Yarats, Ilya Kostrikov, and Rob Fergus. Automatic data augmentation for generalization in deep reinforcement learning, 2021. [1](#)
- [34] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In *European Conference on Computer Vision*, pages 400–418. Springer, 2020. [2](#), [3](#), [6](#), [7](#)
- [35] Santhosh K Ramakrishnan, Dinesh Jayaraman, and Kristen Grauman. An exploration of embodied visual exploration. *International Journal of Computer Vision*, 129(5):1616–1649, 2021. [3](#), [7](#)
- [36] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *International Conference on Learning Representations*, 2018. [1](#), [2](#)
- [37] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [3](#), [6](#)
- [38] Soroush Seifi and Tinne Tuytelaars. Where to look next: Unsupervised active visual exploration on 360° input. *CoRR*, abs/1909.10304, 2019. [1](#)
- [39] Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological slam for visual navigation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12872–12881, 2020. [1](#), [6](#)
- [40] Ivan Sosnovik, Artem Moskalev, and Arnold WM Smeulders. Scale equivariance improves siamese tracking. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2765–2774, 2021. [3](#)
- [41] Ivan Sosnovik, Michał Szmaja, and Arnold Smeulders. Scale-equivariant steerable networks. In *International Conference on Learning Representations*, 2020. [3](#)
- [42] Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. 2021. [1](#)
- [43] Erik Henning Thiede, Truong-Son Hy, and Risi Kondor. The general theory of permutation equivariant neural networks and higher order graph variational encoders. *CoRR*, abs/2004.03990, 2020. [3](#)
- [44] Arnoud Visser, Xingrui-Ji, Merlijn van Ittersum, Luis A. González Jaime, and Laurențiu A. Stancu. Beyond frontier exploration. In Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert, editors, *RoboCup 2007: Robot Soccer World Cup XI*, pages 113–123, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. [1](#)
- [45] Robin Walters, Jinxi Li, and Rose Yu. Trajectory prediction using equivariant continuous convolution. In *International Conference on Learning Representations*, 2021. [3](#)
- [46] Daniel E. Worrall and Max Welling. Deep scale-spaces: Equivariance over scale. In *Advances in Neural Information Processing Systems*, volume 32, pages 7364–7376, 2019. [3](#)
- [47] Yi Wu, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, and Yuandong Tian. Bayesian relational memory for semantic visual navigation. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2769–2779, 2019. [1](#), [2](#)
- [48] Fei Xia, Amir R. Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE, 2018. [2](#), [6](#)
- [49] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, pages 146–151, 1997. [1](#), [2](#), [7](#)
- [50] Denis Yarats, Ilya Kostrikov, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*, 2021. [1](#)
- [51] Joel Ye, Dhruv Batra, Erik Wijmans, and Abhishek Das. Auxiliary tasks speed up learning point goal navigation. In Jens Kober, Fabio Ramos, and Claire Tomlin, editors, *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 498–516. PMLR, 16–18 Nov 2021. [1](#)
- [52] Jaejun Yoo, Namhyuk Ahn, and Kyung-Ah Sohn. Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8375–8384, 2020. [1](#)
- [53] Richard Zhang. Making convolutional networks shift-invariant again. In *International conference on machine learning*, pages 7324–7334. PMLR, 2019. [2](#), [5](#)
- [54] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3357–3364, 2017. [1](#), [2](#)