

# Diagnosing the Environment Bias in Vision-and-Language Navigation

Yubo Zhang\*, Hao Tan\* and Mohit Bansal

UNC Chapel Hill

{zhangyb, haotan, mbansal}@cs.unc.edu

## Abstract

Vision-and-Language Navigation (VLN) requires an agent to follow natural-language instructions, explore the given environments, and reach the desired target locations. These step-by-step navigational instructions are crucial when the agent is navigating new environments about which it has no prior knowledge. Most recent works that study VLN observe a significant performance drop when tested on unseen environments (i.e., environments not used in training), indicating that the neural agent models are highly biased towards training environments. Although this issue is considered as one of the major challenges in VLN research, it is still under-studied and needs a clearer explanation. In this work, we design novel diagnosis experiments via environment re-splitting and feature replacement, looking into possible reasons for this environment bias. We observe that neither the language nor the underlying navigational graph, but the low-level visual appearance conveyed by ResNet features directly affects the agent model and contributes to this environment bias in results. According to this observation, we explore several kinds of semantic representations that contain less low-level visual information, hence the agent learned with these features could be better generalized to unseen testing environments. Without modifying the baseline agent model and its training method, our explored semantic features significantly decrease the performance gaps between seen and unseen on multiple datasets (i.e. R2R, R4R, and CVDN) and achieve competitive unseen results to previous state-of-the-art models.<sup>1</sup>

## 1 Introduction

Vision-and-Language Navigation (VLN) tests an agent’s ability to understand complex natural language instructions, explore the given environments and find the correct paths to the target locations, as shown in Fig. 1. In this work, we focus on

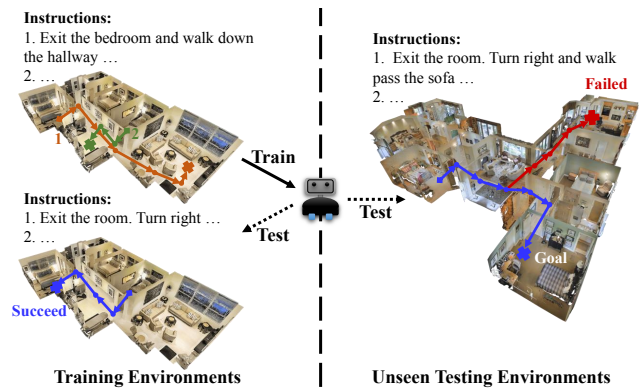


Figure 1: Vision-language-navigation: performance of the agent drops in unseen testing environments.

the instruction-guided navigation [MacMahon *et al.*, 2006; Anderson *et al.*, 2018b; Blukis *et al.*, 2018; Chen *et al.*, 2019b] where detailed step-by-step navigational instructions are used (e.g., ‘Go outside the dining room and turn left ...’), in contrast to the target-oriented navigation [Das *et al.*, 2018; Mirowski *et al.*, 2018; Yu *et al.*, 2019] where only the target is referred (e.g., ‘Go to the kitchen’ or ‘Tell me the color of the bedroom’). Although these step-by-step instructions are over-detailed when navigating local areas (e.g., your home), they are actively used in unseen environments (e.g., a new city) where the desired target is unknown to navigational agents. For this purpose, testing on unseen environments which are not used during agent-training is important and widely accepted by instruction-guided navigation datasets.

Recent works propose different methods to improve the generalizability of agents on these unseen testing environments. Most of the works [Anderson *et al.*, 2018b; Wang *et al.*, 2018; Fried *et al.*, 2018; Wang *et al.*, 2019; Ma *et al.*, 2019a; Ma *et al.*, 2019b; Tan *et al.*, 2019; Huang *et al.*, 2019; Hu *et al.*, 2019] observe a significant performance drop from the environments used in training (seen environments) to the ones not used in training (unseen environments), which indicates the models are strongly biased to the seen environments. While this performance gap is emphasized as one of the major challenges, the issue is still left unresolved and waits for an explicit explanation. Thus, in this paper, we aim to answer three questions to this environment bias: 1. **Where** is the bias

\*Equal Contribution

<sup>1</sup>Code, features at <https://github.com/zhangybzbo/EnvBiasVLN>.

located? 2. **Why** does it exist? 3. **How** to alleviate it?

To locate **where** the bias is, we start by excluding natural-language instructions and underlying navigational graphs from the direct reason for this performance gap. Then in order to conduct a detailed investigation of environments’ effect on agents’ performance, we re-split the environment and categorize the validation data into three sets based on their visibility to the training set: *path-seen*, *path-unseen*, and *env-unseen*. By showing that the results gradually decrease from path-seen to env-unseen data, we thus characterize the ‘spatial localities’ of environment bias at three levels: path level, region level, and environment level. The *low-level information* carried by the ResNet features [He *et al.*, 2016] is our suspect of **why** this locality further leads to the performance gap. By conducting the investigation of replacing the ResNet features with the 1000 ImageNet classification probabilities, we notice that these noisy but higher-level semantic features can effectively reduce the performance gap while maintaining moderate performance on various VLN datasets (i.e., Room-to-Room, R4R, and CVDN), which support our hypothesis.

Following the practice of using the semantic feature in previous investigation experiments, we further provide a discussion on **how** the environment bias could be reduced. Three kinds of more rational and advanced semantic features are adopted: (1) areas of detected objects [Ren *et al.*, 2015]; (2) ground-truth semantic views [Chang *et al.*, 2017]; and (3) learned semantic view features. The results show that all of these semantic features significantly reduce the environment bias in multiple datasets and also achieve strong performance in unseen environments. We hope this work encourages more investigation and research into improving the generalization of vision-language models to unseen real-world scenarios.

## 2 Related Work

**Vision-and-Language Navigation:** Vision-and-language navigation is an emerging area in multi-modality research. Several datasets have been released recently, such as Room-to-Room [Anderson *et al.*, 2018b], Room-for-Room [Jain *et al.*, 2019], Touchdown [Chen *et al.*, 2019b], CVDN [Thomason *et al.*, 2019b] and EQA [Das *et al.*, 2018]. Previous works [Thomason *et al.*, 2019a; Wang *et al.*, 2018; Fried *et al.*, 2018; Wang *et al.*, 2019; Ma *et al.*, 2019a; Ma *et al.*, 2019b; Tan *et al.*, 2019; Hu *et al.*, 2019; Anderson *et al.*, 2019; Ke *et al.*, 2019] focusing on improving the performance of navigation models, especially in unseen testing environments, have helped to increase the navigational success rate.

**Domain Adaptation and Generalization:** The task of domain adaptation is to learn the domain invariant feature with the data samples from different domains. Previous works using adversarial training [Goodfellow *et al.*, 2014; Long *et al.*, 2018] or training transfer functions [Chen *et al.*, 2019a; Rozantsev *et al.*, 2018] have achieved great success. However, in applications where the samples from target domains may not be available (e.g., the testing environments in navigation should not be used in training), these methods are hard to apply. In domain generalization [Blanchard *et al.*, 2011; Muandet *et al.*, 2013; Carlucci *et al.*, 2019], the goal is to

predict the labels in the previous unseen domain, and the testing data is unrevealed during training. In VLN, two previous methods, RCM [Wang *et al.*, 2019] and EnvDrop [Tan *et al.*, 2019], explore the possibility of domain adaptation. Both works take the testing environments in training while RCM also uses testing instructions. In this paper, we focus on the domain generalization problem in VLN, and try to find the reasons for the failures.

## 3 Vision-and-Language Navigation and its Environment Bias

We first introduce the task of Vision-and-Language Navigation (VLN) and briefly describe the neural agent models used in our work. Next, according to the survey of previous works, we show that the environment bias is widely observed in current VLN research.

### 3.1 Vision-and-Language Navigation

**Tasks** As shown in Fig. 1, the agent in VLN task is trained to navigate a certain type of environments  $\{E\}$  given the instruction  $I$ . Each environment  $E$  is an independent space, such as a house, and consists of a set of viewpoints. Each viewpoint is represented as a panoramic image and can be decomposed into separate views  $\{o\}$  as inputs to neural agent models. The viewpoints and their connectivity form the navigational graph. In practice, starting off at a specific viewpoint and provided with an instruction, at each subsequent time step, the agent can observe the panoramic image of the viewpoint where it is located, and choose to move along an edge of the graph to the next node (i.e., viewpoint) or stop. The performance of the agent is evaluated by whether it eventually reaches the target location indicated by the instruction.

**Neural Agent Models** Most instruction-guided navigational agents are built based on attentive encoder-decoder models [Bahdanau *et al.*, 2015]. Based on the instructions which are encoded by the encoder, as well as perceived environments, the decoder generates action outputs. Since the primal focus of this work is to understand the environment bias, we use a standard neural agent model that achieves comparable results to previous works. Specifically, we adopt the panoramic-view neural agent model in [Fried *et al.*, 2018] (‘Follower’) with modifications from [Tan *et al.*, 2019] as our baseline. We also exclude advanced training techniques (i.e., reinforcement learning and data augmentation), only train the agent with imitation learning in all our experiments for the same purpose. Refer to the original papers for more details.

### 3.2 Environment Bias in Indoor Navigation

In order to evaluate the generalizability of agent models, indoor VLN datasets (e.g., those collected from Matterport3D [Chang *et al.*, 2017]) use disjoint sets of environments in training and testing. Two validation splits are provided as well: validation seen (which takes the data from training environments) and validation unseen (whose data is taken from testing environments different from the training).

In the first part of Table 1, we list most of the previous works (R2R [Anderson *et al.*, 2018b], RPA [Wang *et al.*, 2018], S-Follower [Fried *et al.*, 2018], RCM [Wang

Method	Result		
	Val Seen	Val Unseen	Gap $ \Delta $
Room-to-Room [Anderson <i>et al.</i> , 2018b]			
R2R	38.6	21.8	16.8
RPA	42.9	24.6	18.3
S-Follower	66.4	35.5	30.9
RCM	66.7	42.8	23.9
SMNA	67	45	22
Regretful	69	50	19
EnvDrop	62.1	52.2	9.9
ALTR	55.8	46.1	9.7
RN+Obj	59.2	39.5	19.7
CG	31	31	<b>0</b>
Our baseline	56.1	47.5	8.6
<b>Our Learned-Seg</b>	52.6	53.3	<b>0.7</b>
Room-for-Room [Jain <i>et al.</i> , 2019]			
S-Follower	51.9	23.8	28.1
RCM	55.5	28.6	26.9
Our baseline	54.6	30.7	23.9
<b>Our Learned-Seg</b>	38.0	34.3	<b>3.7</b>
CVDN [Thomason <i>et al.</i> , 2019b]			
NDH	5.92	2.10	3.82
Our baseline	6.60	3.05	3.55
<b>Our Learned-Seg</b>	5.82	4.41	<b>1.41</b>
Touchdown [Chen <i>et al.</i> , 2019b]			
GA (original)	7.9 (dev)	5.5 (test)	–
RCONCAT (original)	9.8 (dev)	10.7 (test)	–
Our baseline (original)	15.0 (dev)	14.2 (test)	–
Our baseline (re-split)	17.5	5.3	12.2

Table 1: Results show the performance gaps between seen (‘Val Seen’) and unseen (‘Val Unseen’) environments in several VLN tasks. Room-to-Room and Room-for-Room are evaluated with ‘Success Rate’, CVDN is evaluated with ‘Goal Progress’, Touchdown is evaluated with ‘Task Completion’.

*et al.*, 2019], SMNA [Ma *et al.*, 2019a], Regretful [Ma *et al.*, 2019b], EnvDrop [Tan *et al.*, 2019], ALTR [Huang *et al.*, 2019], RN+Obj [Hu *et al.*, 2019], CG [Anderson *et al.*, 2019]) on the Room-to-Room dataset [Anderson *et al.*, 2018b] and their *success rate* under greedy decoding (i.e., without beam-search) on validation seen and validation unseen splits. The large absolute gaps (from 30.9% to 9.7%) between the results of seen and unseen environments show that current agent models on R2R suffer from environment bias.<sup>2</sup> This phenomenon is also revealed in two other newly-released indoor navigation datasets, Room-for-Room (R4R) [Jain *et al.*, 2019] and Cooperative Vision-and-Dialog Navigation (CVDN) [Thomason *et al.*, 2019b]. The significant result drops from seen to unseen environments can also be observed (i.e., 26.9% on R4R and 3.82 on CVDN), as shown in the second and third parts of Table 1. Lastly,

<sup>2</sup>Our work’s aim is to both close the seen-unseen gap while also achieving competitive unseen results. Note that [Anderson *et al.*, 2019] also achieve 0% gap but at the trade-off of low unseen results.

we show the results (‘Our Learned-Seg’ in Table 1) when the environment bias is effectively reduced by our learned semantic-segmentation features (described in Sec. 6.3), compared to our baselines (denoted as ‘Our baseline’) and previous works.

### 3.3 Environment Bias in Outdoor Navigation

Since the three indoor VLN datasets in previous sections are collected from the same environments, Matterport3D [Chang *et al.*, 2017], in order to demonstrate the generality of this phenomenon, we investigate an outdoor navigation task from Touchdown dataset [Chen *et al.*, 2019b]. In the original data splits of Touchdown, the environment is not specifically divided into seen and unseen. To explore whether the same environment bias can be observed in Touchdown, we split the city environment according to latitude and create two sub-environments: ‘training’ and ‘unseen’. Our baseline neural agent model is adapted with additional convolutional layers to fit this new task. As shown in the last part of Table 1, the big performance gap between the ‘training’ and the ‘unseen’ sub-environment (from 17.5% to 5.3%) still occurs on our re-split data (denoted as ‘re-split’), indicating that environment bias is a broad issue. At the same time, when experimenting on the original data split (denoted as ‘original’), our baseline model achieves state-of-the-art results on the original ‘dev’ set and ‘test’ set, proving the validity of our model in this dataset.

## 4 Where: Effect of Different Task Components

In Sec. 3, we showed that current vision-and-language navigation (VLN) models are biased towards training environments. The purpose of this section is to locate the reason of this environment bias. As one of the early-released and well-explored VLN datasets, Room-to-Room (R2R) [Anderson *et al.*, 2018b] is used as the diagnosing dataset in experiments. We start by showing that two possible candidates, the natural language instructions and the underlying navigational graph, do not directly contribute to the environment bias. Then the effect of visual environments is analyzed in detail.

### 4.1 The Effect of Natural-Language Instructions

A common hypothesis is that the navigational instructions for unseen (e.g., val unseen) and training (i.e., training and val seen) environments are much different due to object and layout differences, and this lingual difference can cause the performance gap. In order to test this hypothesis, we analyze the distributions of success rate with regard to the relationship between the instructions of training and validation sets. To quantitatively evaluate this relationship, we define the ‘distances’ from a validating instruction to all training instructions as the phrase-matching metric. Suppose  $x$  is a validating datum,  $\text{inst}(x)$  is the instruction of this datum, and  $\mathbb{T}$  is the training set, we use ROUGE-L [Lin, 2004] and BLEU-4 [Papineni *et al.*, 2002] to calculate this ‘distance’:

$$\text{dis}_{\text{ROUGE}}(x, \mathbb{T}) = \min_{t \in \mathbb{T}} \text{ROUGE-L}(\text{inst}(x), \text{inst}(t)) \quad (1)$$

$$\text{dis}_{\text{BLEU}}(x, \mathbb{T}) = \text{BLEU-4}(\text{inst}(x), \{\text{inst}(t)\}_{t \in \mathbb{T}}) \quad (2)$$

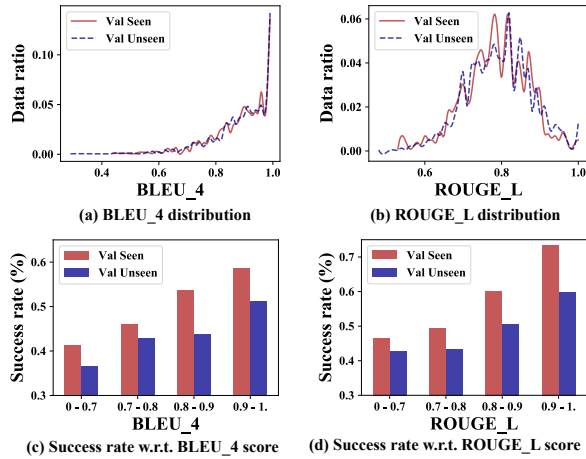


Figure 2: The language 'distance' distribution (defined by language scores) and its relationship to success rate.

where we consider all the training instructions as references in calculating the BLEU-4 score.

The distributions of success rates and distances are shown in Fig. 2. As opposed to the hypothesis, we do not observe a significant difference between the 'distances' distributions (as shown in Fig. 2 (a, b)) of seen and unseen validation data. As for the success rate distributions (in Fig. 2(c,d)), the performance is better when the instruction has smaller 'distances' (i.e., higher BLEU-4/ROUGE-L scores w.r.t. the training instructions) on both validation splits. However, comparing two splits, with the same 'distance' to training instructions, seen validation still significantly outperforms the unseen validation set on success rate, which excludes the language from the cause of this performance gap.

### 4.2 The Effect of Underlying Navigational Graph

An environment could be considered as its underlying navigational graph (as in Fig. 3) with additional visual information (as in Fig. 1). In order to test whether the agent model could overfit to these navigational graphs (and thus be biased towards training environments), we follow the experiments in [Hu *et al.*, 2019] to train the agent without visual information, masking out the ResNet features which are used as the visual input with zero vectors. When the agent could only make the decision based on the instruction and the navigational graph, our baseline model gives the success rate of 38.5% and 41.0% on validation seen and unseen in this setting, which is consistent with the finding in [Hu *et al.*, 2019]. Besides showing the relatively good performance on unseen split without visual contents (similar to [Thomason *et al.*, 2019a] and [Hu *et al.*, 2019]), the model generates the low performance gap between seen and unseen environments (2.5% compared to the > 10% gap in usual). Hence, we claim that the underlying graph is not a dominant reason for the environment bias.

### 4.3 The Effect of Visual Environments

To show how the visual information plays its role in environment bias, we analyze agent's performance on unseen environments and in different spatial regions of the training envi-

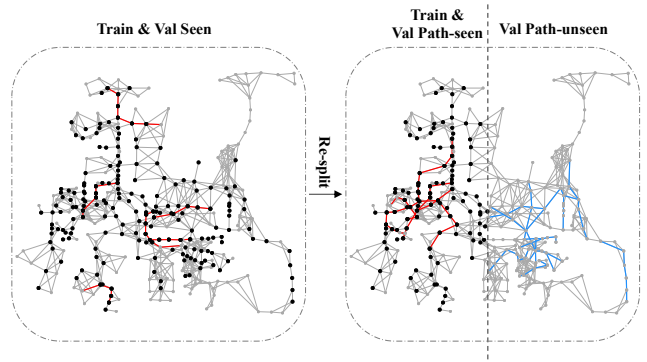


Figure 3: Graph split: left is original data and right is re-splitting data. Black vertices are viewpoints visited during training; red paths are val seen (in the left part) / val path-seen (in the right part); blue paths are val path-unseen.

ronments. In order to give a detailed characterization of the effect of environments, we will reveal three different levels of spatial localities which are related to the agent's performance:

- **Path-level Locality:** Agents are better at paths which intersect with the training paths.
- **Region-level Locality:** Agents are better in regions that are closer to the training data.
- **Environment-level Locality:** Agents perform better in training environments than in unseen environments.

The existence of these spatial locality inspires us to find the direct cause of the problem in Sec. 3.2. However, the original split of data is not fine-grained enough to separately reveal these spatial localities. As we visualize the data from one environment of the Room-to-Room dataset in the left part of Fig. 3, where the vertices are viewpoints with visual information and edges are valid connections between viewpoints, nearly all viewpoints in val-seen paths (vertices connected to red lines) are also included in training paths (vertices marked by dark-black). We thus cannot categorize the path-level and region-level localities. To bypass this, we propose a novel re-splitting method to create our diagnosis data splits.

**Structural Data Re-splitting** We employ two kinds of structural data splitting methods based on the horizontal or vertical coordinates, denoted as 'X-split' and 'Z-split' respectively. The 'Z-split' separates different floors of houses and 'X-split' creates separate areas. When applying to training environments in R2R dataset, we use one side of the splitting line (see 'X-splitting line' in Fig. 3) as the new training 'environment', the other side as the path-unseen 'environment'. The original training and val-seen data are re-split accordingly, while the val-unseen data and environments keep untouched. As shown in the right part of Fig. 3, our re-splitting method creates three new data splits: training split, val-path-seen split (i.e., data intersecting with training paths), and val-path-unseen split (i.e., data not intersecting with training paths but located in the same original environments). The original val-unseen set is denoted as 'env-unseen' split.

**Existence of Path-level and Environment-level Localities** For both splitting methods, we train our baseline model on the

	Splitting	Train	Validation		
			P-seen	P-unseen	E-unseen
Env	R2R	61	56	0	11
	X-split	61	57	16	11
	Z-split	61	56	29	11
Data	R2R	14,025	1,020	0	2,349
	X-split	11,631	1,230	1,098	2,349
	Z-split	10,894	867	2,324	2,349
SR	R2R	88.3	56.1	–	47.5
	X-split	87.3	58.9	52.6	46.7
	Z-split	94.7	62.5	47.8	42.4

Table 2: The statistics and results of original R2R splits and our new splits, showing the path-level and environment-level localities. Numbers of environments and data are denoted as ‘Env’ and ‘Data’; success rate is denoted as ‘SR’; val-path-seen, val-path-unseen and val-env-unseen are denoted as ‘P-seen’, ‘P-unseen’ and ‘E-unseen’.

newly-split training set and evaluate it on the three validation sets. As shown in Table 2, the agent performs better on val path-seen than val path-unseen, which suggests that a **path-level locality** exists in current VLN agent models. Meanwhile, the results on val path-unseen are then higher than val env-unseen, indicating the existence of **environment-level locality** which is independent of the path-level locality.

**Existence of Region-level Locality** To demonstrate region-level locality, we study how the success rate changes with respect to distances between training and validation environment regions, similar to the analysis of language ‘distance’ in Sec. 4.1. We calculate the point-by-point shortest paths using the Dijkstra’s algorithm, where the shortest distances between viewpoints  $v$  and  $v'$  are denoted as the graph distance  $\text{dis}_{\text{GRAPH}}(v, v')$ . We first define the viewpoint distance  $\text{dis}_{\text{VIEWPOINT}}$  from a viewpoint  $v$  to the training data  $\mathbb{T}$  as the minimal graph distance from  $v$  to a viewpoint  $v'$  in training data; Then the path distance  $\text{dis}_{\text{PATH}}$  from a validating data  $x$  to the whole training data  $\mathbb{T}$  is defined as the maximal viewpoint distance of the viewpoints in the path of  $x$ :

$$\text{dis}_{\text{PATH}}(x, \mathbb{T}) = \max_{v \in \text{path}(x)} \text{dis}_{\text{VIEWPOINT}}(v, \mathbb{T}) \quad (3)$$

$$\text{dis}_{\text{VIEWPOINT}}(v, \mathbb{T}) = \min_{v' \in \text{path}(t)} \text{dis}_{\text{GRAPH}}(v, v') \quad (4)$$

We compute path distances between paths in the path-unseen validation set and training environments of our re-splitting data. As shown in Table 3, the success rate declines as the path moves further from the training environment in both re-splitting methods (i.e., ‘X-split’ and ‘Z-split’). Therefore, the closer the path to the training data, the better the agent performs, which reveals the existence of **region-level locality**.

## 5 Why: What Inside the Environments Contributes to the Bias?

In Sec. 4, we locate the cause of the performance gap in visual environments by excluding other potential reasons and categorizing the spatial localities. However, there are still multiple possible aspects inside the environment which could lead

X-split	PD (meters)	5-13	14-16	17-21	22-57
	SR (%)	56.3	56.2	50.5	43.6
Z-split	PD (meters)	5-10	11-13	14-17	18-52
	SR (%)	58.5	47.9	42.9	44.1

Table 3: The success rate declines as the path moves further from training regions. Path distance is denoted as ‘PD’.

to these spatial localities, e.g., the objects layout and room connections. The agent could be biased towards the training environments by over-memorizing these environment-specific characteristics. In this section, we pinpoint the aspect that directly contributes to the bias to be the low-level visual information carried by the ResNet features. We first show an experiment where the gap between seen and unseen environments is effectively decreased by minimal model modifications. We then clarify our conclusions based on the findings.

### 5.1 An Investigation Experiment: ImageNet Labels as Visual Features

Suspecting that the over-fitting is caused by low-level ResNet 2048-features which the agent over-learns, we hope to find a replacement that contains minimal low-level information while preserving distinguishable visual contents. The most straightforward replacement is that instead of using mean-pooled features, we inject the frozen 1000-way classifying layer from ResNet pre-training, and use the probabilities of ImageNet labels as visual features. Shown as ‘ImageNet’ in Table 4, the probability distribution reduces the gaps between seen and unseen in all three datasets, and almost closes the gaps of R2R and R4R. These results further reveal the low-level ResNet features of image views as the reason for environment bias. Combining with the findings of spatial localities, we suggest that environments (i.e., houses) and regions (i.e., rooms) usually have their own ‘styles’. Thus the same semantic label (captured by ImageNet-1000 features) may have different visual appearances (captured by ResNet features) in different environments or regions. As a result, ImageNet-1000 features, in spite of being noisy, are not distracted by low-level visual appearance and thus could generalize well to unseen environments, while ResNet features could not.

Although this ImageNet-1000-feature replacement could decrease the performance gap, it has a disadvantage with the VLN domain that the validation unseen result of R4R is slightly worse than baseline (and not much better in R2R and CVDN cases). Hence it motivates us to find better semantic representations of environmental features that can both close the seen-unseen gap while also achieving state-of-the-art on unseen results (which we will discuss next).

## 6 How: Methodology to Fix Environment Bias

In the previous section (Sec. 5), we found that the environment bias is related to the low-level visual information (i.e., 2048-dim ResNet features), and we want to build our agent on the features which are correlated to the VLN environmental semantics following the observation in Sec. 5.1. In this

Task	Feature	Result		
		Val Seen	Val Unseen	Gap $ \Delta $
R2R	ResNet	56.1	47.5	8.6
	ImageNet	47.1	48.2	1.1
	Detection	55.9	50.0	5.9
	GT-Seg	55.6	56.2	0.6
	Learned-Seg	52.6	53.3	0.7
R4R	ResNet	54.6	30.7	23.9
	ImageNet	28.7	28.9	0.2
	Detection	48.8	32.0	16.8
	GT-Seg	47.6	35.9	11.7
	Learned-Seg	38.0	34.3	3.7
CVDN	ResNet	6.60	3.05	3.55
	ImageNet	5.70	3.11	2.59
	Detection	6.55	3.94	2.61
	GT-Seg	6.46	4.33	2.13
	Learned-Seg	5.82	4.41	1.41

Table 4: Results showing that our semantic feature representations, i.e., ImageNet, Detection, ground-truth and learned semantic segmentation (denoted as ‘GT-Seg’ and ‘Learned-Seg’), effectively reduce the performance gaps in all three datasets.

section, several advanced semantic feature replacements are explored. As shown in Table 4, these features could effectively reduce the performance gaps between seen and unseen environments and improve the unseen results compared to our strong baselines, without any changes in training procedure or hyperparameters. The effectiveness of these semantic representations supports our explanation of the environment bias in Sec. 5 and also suggests that future work in VLN tasks should think about such generalization issues.

### 6.1 Baseline

In our baseline model (described in Sec. 3.1), following the previous works we use the standard ResNet features as the representation of environments [Anderson *et al.*, 2018b; Jain *et al.*, 2019; Thomason *et al.*, 2019b]. These features come from the mean-pooled layer after the final convolutional layer of ResNet-152 [He *et al.*, 2016] pre-trained on ImageNet [Russakovsky *et al.*, 2015]. As shown in ‘Baseline’ rows of Table 4, in this setting, val-seen results are significantly higher than val-unseen results in all three datasets. The ‘environment bias’ phenomenon, which we described in Sec. 3.2, is observed in our baseline model, leading us to the following discussions of semantic features.

### 6.2 Detected Object Areas

During navigation, the objects in the environments are crucial since their matchings with instruction often indicate the locations that can guide the agent, thus object detection quality of the environments can provide relevant semantic information. In our work, we utilize the detection information generated by Faster R-CNN [Ren *et al.*, 2015] to create the feature representations. Compared to ImageNet-1000 features (Sec. 5.1), these detection features include more environmental information since the viewing images in VLN usually contain multiple objects. Instead of directly using classification probabil-

ities or the label embeddings (in [Hu *et al.*, 2019]), we design our features  $f_{\text{DETECT}}$  of each image view as the collection of the area summations weighted by detection confidence of each detected object:

$$f_{\text{DETECT}}=[a_{c_1}, a_{c_2}, \dots, a_{c_n}]; \quad (5)$$

$$a_{c_i} = \sum_{\text{obj is } c_i} \text{Area}(\text{obj}) \cdot \text{Conf}(\text{obj}) \quad (6)$$

where  $a_{c_i}$  is feature of the object label  $c_i$ . The area and confidence are denoted as  $\text{Area}(\ast)$  and  $\text{Conf}(\ast)$ . For implementation details, we use the Faster R-CNN [Ren *et al.*, 2015] trained on Visual Genome [Krishna *et al.*, 2017] provided in Bottom-Up Attention [Anderson *et al.*, 2018a]. To eliminate labels irrelevant to the VLN task, we calculate the total areas of each detection object among all environments and pick the labels that take up a relatively large proportion of the environments, creating features of dimension 152. The results are denoted as ‘Detection’ in Table 4. The performance gaps are diminished with these detection features compared to baselines in all three datasets, indicating that lifting the features to a higher semantic level has a positive effect on alleviating the environment bias. Meanwhile, the improvement of unseen validation results suggests better effectiveness in the VLN task than the ImageNet labels.

### 6.3 Semantic Segmentation

Although the detection features can provide adequate semantic information for the agent to achieve comparable results as the baseline model, they do not fully utilize the visual information where the content left over from detection may contain useful knowledge for navigation. A better semantic representation is the semantic segmentation, which segments each view image on the pixel level and gives the label to each segment region. Matterport3D [Chang *et al.*, 2017] dataset provides the labeled semantic segmentation information of every scene and we take the rendered images from [Tan *et al.*, 2019]. An example of comparison between RGB images and semantic segmentation is available in the Appendix. Since the semantic segmentation images are noisy and blurry in boundaries, following the design of detection features, the semantic features are designed as the areas of semantic classes in each image view, with dimension of 42. We first assume that the semantic segmentation is provided as additional environmental information and the results of the model using the ground-truth semantic segmentation areas are shown in the ‘GT-Seg’ rows in Table 4. We next study the situation where the semantic information is not available in testing environments thus the information needs to be learned from training environments. A separate multi-layer perceptron is trained to predict the areas of these semantic classes from 2048-dim ResNet features. This MLP is trained on 51 (out of 61) training environments and its hyper-parameters are validated on the remaining training 10 environments (and more details are available in the Appendix). The results of the agent model with these predicted semantics as features are shown in ‘Learned-Seg’ As shown in Table 4, both ‘GT-Seg’ and ‘Learned-Seg’ semantic representations bring the performance of seen and unseen even closer compared to the previ-



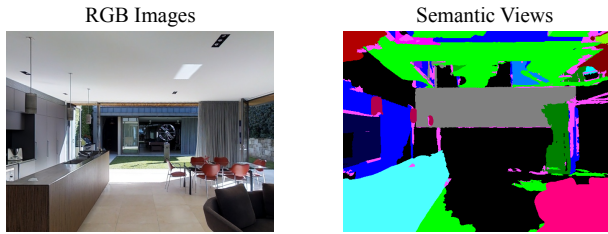


Figure 4: Comparisons between RGB images and their semantic views.

ous detection semantics. The highest validation unseen success rates among our proposed representations are also produced by these methods: ‘Learned-Seg’ semantics for CVDN and ‘GT-Seg’ semantics for R2R and R4R. Overall, among all explored visual representations, the semantic segmentation features are most effective in reducing the environment bias.<sup>3</sup>

## 7 Conclusion

In this paper, we focus on studying the performance gap between seen and unseen environments widely observed in vision-and-language navigation (VLN) tasks, trying to find where and why this environment bias exists and provide possible initial solutions. By designing the diagnosis experiments of environment re-splitting and feature replacement, we locate the environment bias to be in the low-level visual appearance; and we discuss semantic features that decrease the performance gaps in three VLN datasets and achieve state-of-the-art results.

## Acknowledgments

We thank P. Hase, H. Guo, L. Yu, S. Zhang, and X. Zhou for their helpful suggestions. This work was supported by ARO-YIP Award W911NF-18-1-0336, ONR Grant N00014-18-1-2871, NIH grant 1R01LM013329-01, and a Google Focused Research Award. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the funding agency.

## A Example of RGB Images and Semantic Segmentation

In Fig. 4, we show a rendered semantic view from [Tan *et al.*, 2019] and its original RGB image. Different colors indicate the label of segmentation areas. 42 semantic labels (40 semantic classes and 2 structural classes) are considered in the Matterport3D dataset [Chang *et al.*, 2017].

## B Details of ‘Learned-Seg’ Semantic Training

We use a multi-layer perceptron (MLP) to generate the ‘Learned-Seg’ semantic features. The multi-layer perceptron

<sup>3</sup>To further prove the effectiveness of our semantic features in alleviating the environment bias, we test our ‘Learned-Seg’ features on two other methods from previous works (SMNA and Regretful). See details in Appendix.

Models	Features	Result		
		Val Seen	Val Unseen	Gap $ \Delta $
SMNA	ResNet	63.9	38.0	25.9
	Learned-Seg	55.0	40.6	14.4
Regretful	ResNet	63.8	42.7	21.1
	Learned-Seg	55.2	47.3	7.9

Table 5: When being used in two other agent models, our learned semantic segmentation features still show the effectiveness in alleviating the environment bias.

includes three fully-connected layers with ReLU activation on the outputs of the first two layers. The input of this MLP is the 2048-dim ResNet feature  $f$  of each image view. The hidden sizes of the first two layers are 512 and 256. The final layer will output the 42-dim semantic feature  $y$  that represents the areas of each semantic class. After the linear layers, we use the sigmoid function  $\sigma$  to convert the output to the ratio of areas.

$$x_1 = \text{ReLU}(A_1 f + b_1) \quad (7)$$

$$x_2 = \text{ReLU}(A_2 x_1 + b_2) \quad (8)$$

$$y = \sigma(A_3 x_2 + b_3) \quad (9)$$

The model is trained with ground-truth semantic areas (normalized to  $[0, 1]$ ) of the views in 51 environments out of total 61 VLN training environments, and is tuned on the remaining 10 environments. We minimize the binary cross-entropy loss between the ground-truth areas  $\{y_i^*\}$  and the predicted areas  $\{y_i\}$ , where  $i$  indicate the  $i$ -th semantic class.

$$\mathcal{L} = - \sum_i (y_i^* \log y_i + (1 - y_i^*) \log (1 - y_i)) \quad (10)$$

Dropout layers with a probability of 0.3 are added between fully-connected layers while training.

After the model is fitted, we freeze the weight and use it to predict the ‘Learned-Seg’ semantic features of all seen and unseen environments (i.e., environments of all training, val-seen, and val-unseen data). These extracted features are further used in training the neural navigational agent model.

## C Other Navigational Models with ‘Learned-Seg’ Features

To further prove the effectiveness of our semantic features in alleviating the environment bias, we test our ‘Learned-Seg’ features (described in Sec. 6.3) on two other baselines from previous works, i.e., SMNA [Ma *et al.*, 2019a] and Regretful [Ma *et al.*, 2019b].<sup>4</sup> As shown in Table 5, when being used in different neural agent models, our semantic features can also effectively reduce the performance gaps as well as bring extra improvement on val unseen results.

<sup>4</sup>We use their official code implementation in <https://github.com/chihaoma/selfmonitoring-agent> and <https://github.com/chihaoma/regretful-agent>

## References

- [Anderson *et al.*, 2018a] P. Anderson, X. He, et al. Bottom-up and top-down attention for image captioning and visual question answering. In *CVPR*, 2018.
- [Anderson *et al.*, 2018b] P. Anderson, Q. Wu, et al. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *CVPR*, 2018.
- [Anderson *et al.*, 2019] P. Anderson, A. Shrivastava, et al. Chasing ghosts: Instruction following as bayesian state tracking. In *NeurIPS*, 2019.
- [Bahdanau *et al.*, 2015] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- [Blanchard *et al.*, 2011] G. Blanchard, G. Lee, and C. Scott. Generalizing from several related classification tasks to a new unlabeled sample. In *NeurIPS*, 2011.
- [Blukis *et al.*, 2018] V. Blukis, D. Misra, et al. Mapping navigation instructions to continuous control actions with position-visitation prediction. In *CoRL*, 2018.
- [Carlucci *et al.*, 2019] F. M. Carlucci, A. D’Innocente, et al. Domain generalization by solving jigsaw puzzles. In *CVPR*, 2019.
- [Chang *et al.*, 2017] A. Chang, A. Dai, et al. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, 2017.
- [Chen *et al.*, 2019a] C. Chen, Z. Chen, et al. Joint domain alignment and discriminative feature learning for unsupervised deep domain adaptation. In *AAAI*, 2019.
- [Chen *et al.*, 2019b] H. Chen, A. Suhr, et al. Touchdown: Natural language navigation and spatial reasoning in visual street environments. In *CVPR*, 2019.
- [Das *et al.*, 2018] A. Das, S. Datta, et al. Embodied question answering. In *CVPR*, 2018.
- [Fried *et al.*, 2018] D. Fried, R. Hu, et al. Speaker-follower models for vision-and-language navigation. In *NeurIPS*, 2018.
- [Goodfellow *et al.*, 2014] I. Goodfellow, J. Pouget-Abadie, et al. Generative adversarial nets. In *NeurIPS*, 2014.
- [He *et al.*, 2016] K. He, X. Zhang, et al. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Hu *et al.*, 2019] R. Hu, D. Fried, et al. Are you looking? grounding to multiple modalities in vision-and-language navigation. In *ACL*, 2019.
- [Huang *et al.*, 2019] H. Huang, V. Jain, et al. Transferable representation learning in vision-and-language navigation. In *ICCV*, 2019.
- [Jain *et al.*, 2019] V. Jain, G. Magalhaes, et al. Stay on the path: Instruction fidelity in vision-and-language navigation. In *ACL*, 2019.
- [Ke *et al.*, 2019] L. Ke, X. Li, et al. Tactical rewind: Self-correction via backtracking in vision-and-language navigation. In *CVPR*, 2019.
- [Krishna *et al.*, 2017] R. Krishna, Y. Zhu, et al. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *IJCV*, 123(1):32–73, 2017.
- [Lin, 2004] C. Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [Long *et al.*, 2018] M. Long, Z. Cao, et al. Conditional adversarial domain adaptation. In *NeurIPS*, 2018.
- [Ma *et al.*, 2019a] C. Ma, J. Lu, et al. Self-monitoring navigation agent via auxiliary progress estimation. In *ICLR*, 2019.
- [Ma *et al.*, 2019b] C. Ma, Z. Wu, et al. The regretful agent: Heuristic-aided navigation through progress estimation. In *CVPR*, 2019.
- [MacMahon *et al.*, 2006] M. MacMahon, B. Stankiewicz, and B. Kuipers. Walk the talk: connecting language, knowledge, and action in route instructions. In *AAAI*, 2006.
- [Mirowski *et al.*, 2018] P. Mirowski, M. Grimes, et al. Learning to navigate in cities without a map. In *NeurIPS*, 2018.
- [Muandet *et al.*, 2013] K. Muandet, D. Balduzzi, and B. Schölkopf. Domain generalization via invariant feature representation. In *ICML*, 2013.
- [Papineni *et al.*, 2002] K. Papineni, S. Roukos, et al. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002.
- [Ren *et al.*, 2015] S. Ren, K. He, et al. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [Rozantsev *et al.*, 2018] A. Rozantsev, M. Salzmann, and P. Fua. Residual parameter transfer for deep domain adaptation. In *CVPR*, 2018.
- [Russakovsky *et al.*, 2015] O. Russakovsky, J. Deng, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [Tan *et al.*, 2019] H. Tan, L. Yu, and M. Bansal. Learning to navigate unseen environments: Back translation with environmental dropout. In *NACCL*, 2019.
- [Thomason *et al.*, 2019a] J. Thomason, D. Gordan, and Y. Bisk. Shifting the baseline: Single modality performance on visual navigation & qa. In *NACCL*, 2019.
- [Thomason *et al.*, 2019b] J. Thomason, M. Murray, et al. Vision-and-dialog navigation. In *CoRL*, 2019.
- [Wang *et al.*, 2018] X. Wang, W. Xiong, et al. Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In *ECCV*, 2018.
- [Wang *et al.*, 2019] X. Wang, Q. Huang, et al. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *CVPR*, 2019.
- [Yu *et al.*, 2019] L. Yu, X. Chen, et al. Multi-target embodied question answering. In *CVPR*, 2019.